



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

First-Passage Kinetic Monte Carlo. II: An Algorithm for Complex Diffusion-Reaction Systems

A. Donev, V. Bulatov, T. Oppelstrup, G. Gilmer, B. Sadigh, M. Kalos

May 21, 2009

Journal of Computational Physics

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

First-Passage Kinetic Monte Carlo. II: An Algorithm for Complex Diffusion-Reaction Systems

Aleksandar Donev,¹ Vasily V. Bulatov,¹ Tomas Oppelstrup,^{1,2}

George H. Gilmer,¹ Babak Sadigh,¹ and Malvin H. Kalos¹

¹*Lawrence Livermore National Laboratory, Livermore, California 94551, USA*

²*Royal Institute of Technology (KTH), Stockholm S-10044, Sweden*

Abstract

In this second part of a series of two papers we develop an asynchronous event-driven First-Passage Kinetic Monte Carlo (FPKMC) algorithm for continuous time and space systems involving multiple diffusing and reacting species of spherical particles in two and three dimensions. The FPKMC algorithm presented here is a generalization of the method described in the first part of this series of papers and is implemented in a robust and flexible framework. Unlike standard KMC algorithms such as the n -fold algorithm, FPKMC is most efficient at low densities where it replaces the many small hops needed for reactants to find each other with large first-passage hops sampled from exact time-dependent Green's functions, *without* sacrificing accuracy. We describe in detail the key components of the algorithm, including the event-loop and the sampling of first-passage probability distributions, and demonstrate the accuracy of the new method. We apply the FPKMC algorithm to the challenging problem of simulation of long-term irradiation of metals, relevant to the performance and aging of nuclear materials in current and future nuclear power plants. The problem of radiation damage spans many decades of time-scales, from picosecond spikes caused by primary cascades, to years of slow damage annealing and microstructure evolution. Our implementation of the FPKMC algorithm has been able to simulate the irradiation of a metal sample for durations that are orders of magnitude longer than any previous simulations using the standard Object KMC or more recent asynchronous algorithms.

I. INTRODUCTION

In the first part of this series of two papers we described a novel diffusion Kinetic Monte Carlo (KMC) algorithm for simulating diffusion-reaction systems in one, two, and three dimensions. The algorithm cures the notorious inefficiency of standard KMC algorithms [1] at low densities of the reacting and diffusing particles: when diffusion is simulated via a sequence of small hops, many such hops are needed to bring reactants together and the standard algorithms are unable to progress over sufficiently long time scales. First proposed in Ref. [2] and described in detail in the first part of this series of papers, the essential idea behind the new First Passage Kinetic Monte Carlo (FPKMC) algorithm is to replace the long sequences of small hops with large super-hops sampled from an exact Green's function derived for a simpler auxiliary problem in which the diffusing particles are spatially isolated (protected) from each other and thus diffuse independently. The resulting algorithm is not only fast even at low densities, but it is also exact (modulo numerical precision) and does not, in principle, introduce approximations employed in other accelerated algorithms, such as discretizing continuum diffusion into a sequence of hops [1] or neglecting less-probable reactions [3, 4].

In the first part of this series of two papers we focused on the basic ideas behind the new FPKMC method and illustrated its application in the context of simple models in which identical or nearly identical particles randomly walk on a lattice or in the continuum. In this second part we extend the algorithm to more general and complex diffusion-reaction systems, such as multiple species with different sizes and diffusion coefficients, competing reaction mechanisms, e.g. particle conversion, death and insertion, absorbing boundary conditions, focusing on the case of continuum isotropic diffusion without advection. We will give a general and formal presentation of the FPKMC method as an event-driven asynchronous algorithm implemented for the case of additive hard spheres in two and three dimensions. We will present algorithmic details, including detailed pseudo-codes, for a flexible yet efficient implementation of the FPKMC method capable of handling a variety of problems of interest, including radiation damage in metals [3, 5, 6], dopant implantation [1], surface reactions [7, 8], coarsening [9], (bio)chemical reaction networks [4, 10, 11], and others. We validate the new algorithm by comparing our simulation results against the Object KMC BIGMAC code [1] for several non-trivial test problems. We then apply the new algorithm

to simulations of irradiation damage accumulation in iron and validate the new method by comparing to results obtained using the existing object KMC code LAKIMOCA [6]. Finally, we demonstrate that the new FPKMC algorithm allows to extend the time horizon of radiation damage simulations well beyond current computational limits and to reach, for the first time, the long time-scales of material life in a nuclear reactor (section V B 3).

In the remainder of this section we specify our diffusion-reaction model and briefly discuss asynchronous event-driven algorithms. In Section II we describe the core of the FPKMC algorithm, namely, the use of exact time-dependent Green's functions for a suitably-defined separable sub-problem. The new method is rather general and extends to a variety of problems where diffusion plays a role, including also discrete (lattice) systems and more general types of diffusion. In Section III we discuss numerical evaluation of the time-dependent Green's functions for the case of hard spheres. Section IV gives further details, including detailed pseudocodes for key components of the FPKMC algorithm. In Section V we present numerical validation of the algorithm along with some performance figures, and finally, in Section VI we offer a few concluding remarks.

A. Model Representation

Consider a simulation of the time evolution of a collection of N diffusing reactive particles in d -dimensions. For simplicity, we will focus on the case of hard spheres of fixed radius diffusing in a homogeneous medium. In the absence of reactions or surfaces and assuming the particle started in some specific point \mathbf{r}_0 at time 0, the probability $c(\mathbf{r}, t)$ of finding the same particle in position \mathbf{r} at time t is the solution to the time-dependent diffusion equation

$$\partial_t c = D \nabla^2 c \text{ and } c(\mathbf{r}, 0) = \delta(\mathbf{r} - \mathbf{r}_0), \quad (1)$$

where D is the particle diffusion coefficient.

At any point in time, the state of the system is characterized by its *configuration* $\mathbf{Q} = (\mathbf{q}_1, \dots, \mathbf{q}_N)$. The number of particles N may itself vary with time. Each particle i can possess an arbitrary number of attributes \mathbf{a}_i in addition to the position of its centroid \mathbf{r}_i , $\mathbf{q}_i = (\mathbf{r}_i, \mathbf{a}_i)$. These attributes include a species $1 \leq s_i \leq N_s$, radius R_i , diffusion coefficient D_i , as well as other problem-specific attributes such as charge, mass, etc. Some attributes may be shared by all particles of a given species, for example, all particles of a given species

may have the same diffusion coefficient.

The symmetric *reaction table* $\mathcal{R}_{\alpha\beta}$ of size $N_s(N_s + 1)/2$ specifies the type of *two-particle reaction* that occurs when particles A and B of species α and β collide. In particular, $\mathcal{R}_{\alpha\beta}$ can specify that particles of species α and β do not interact with each other. Example of possible reactions are *annihilation* $A + B \rightarrow 0$, *chemical reaction* $A + B \rightarrow C$, including the special case of *absorption* $A + B \rightarrow A$, *coalescence* $A + B \rightarrow AB$, and *reflection* $A + B \rightarrow A + B$. The *decay table* \mathcal{D}_k^α specifies an arbitrary number N_d^α of possible *single-particle reactions* for particles of species α , assumed to occur as a Poisson process with rates $\Gamma_k^\alpha = (\tau_k^\alpha)^{-1}$. Examples of decay reactions are *splitting* $A \rightarrow B + C$, including the special case of *emission* $A \rightarrow A + B$, *transmutation* $A \rightarrow B$, *death* $A \rightarrow 0$, and *jump* (kick) $A \rightarrow A$. The positions of any products of the single- or two-particle reactions are assigned depending on the positions and attributes of the reactants, sometimes with additional random displacements. If hard-wall boundaries are present, particles colliding with a *hard wall* k they may be absorbed or reflected with certain pre-specified probabilities.

The *insertion rates* \mathcal{B}_α specify the rate of insertion (birth) for particles of species α per unit time per unit volume. Typically particles are inserted randomly and uniformly inside the simulation volume.

B. Asynchronous Event-Driven Algorithms

The First-Passage Kinetic Monte Carlo (FPKMC) algorithm belongs to the category of *asynchronous* event-driven (AED) algorithms [12]. The algorithm is similar to the well-known event-driven molecular dynamics (EDMD) algorithm [13], with the essential difference that in FPKMC particle dynamics is stochastic rather than deterministic [14]. Just like the hard-sphere Molecular Dynamics, FPKMC is an *exact* algorithm (within numerical precision) because the two-particle problem in FPKMC can be solved exactly. By its exactness we mean that FPKMC samples evolution trajectories of N random walkers from the correct probability distribution, as given by the exact solution of an appropriate Master equation.

Event-driven algorithms evolve the state of the system by updating it only when certain non-trivial *events* occur, skipping the time elapsed between such events as uninteresting, e.g., unchanged or analytically solvable. In the asynchronous algorithms, there is a global simulation time t , typically the time when the last processed event occurred, and each

particle i is associated with a point in time $t_i \leq t$, typically the last time it participated in an event. This is to be contrasted to synchronous event-driven algorithms, where all of the particles are at the same time t , such as the n -fold (BKL) algorithm for performing kinetic (dynamic) Monte Carlo simulations [15]. The classical n -fold algorithm hinges on the fact that the state of the system does not change between the events, as is common in lattice models where particle positions are discrete. For example, the atoms may vibrate around the lattice sites and occasionally hop to nearby sites. In the model considered here, however, the positions of the particles are continuous and continuously changing even between events.

An asynchronous simulation progresses by processing the event at the head of the queue, scheduling new events for any affected particles, and then updating the event queue. This generic procedure is summarized in Algorithm 1. The main types of events in FPKMC and their scheduling and processing will be described in detail in the next section.

Each particle i stores the time it was last updated t_i along with a prediction for its *impending event* (t_i^e, p_i, ν_i) , specified via the predicted time of occurrence (timestamp) t_i^e , the event *partner* p_i , and the event *qualifier* (type of event) ν_i . When it is clear what particle we are referring to, we will omit the subscript i for simplicity. If several different events are possible for a given particle then the first event scheduled to occur (i.e., the one with the smallest t_e) is chosen. The event times for all particles and event times for any *external events* are stored in a priority queue (e.g., a heap) called the *event queue*. If the partner $p = j$ is another particle, then event prediction for the partner particle j is not stored in the event queue to avoid sorting duplicate events with equal timestamps.

Algorithm 1: The event loop in the asynchronous event-driven algorithms.

1. *De-queue* the event at the top of the event queue, of type ν and timestamp t_e , to find the next particle i to have an event, or the next external event to occur.
2. *Advance* the global simulation time $t \leftarrow t_e$, and update the configuration of all involved particles to the current simulation time (e.g., $t_i \leftarrow t_e$).
3. *Process* the event. If necessary update the other particles (e.g., if p is another particle) to the current simulation time and re-schedule their events. If an external event or particle i no longer exists, cycle to the next event in the queue.

4. *Schedule* a new event for particle i at time t_e , and insert it into the event queue with the new timestamp (key) $t_i \leftarrow t_e$.
 5. *Cycle* back to step 1.
-

II. FIRST-PASSAGE KINETIC MONTE CARLO ALGORITHM

A detailed description of the FPKMC algorithm, including pseudo-codes, is given in Section IV. Here we only briefly discuss the most important components of the algorithm. Although this description is intended to be self-contained, the reader is referred to the first part of this series of papers for a more intuitive introduction.

The essential idea behind the FPKMC algorithm is to break the N -body problem into a collection of independent one-body or two-body (pair) problems that can be solved analytically. This is achieved by protecting each particle i with a *protective region* \mathcal{P}_i , $\mathcal{C}_i \subseteq \mathcal{P}_i$, where the hard core of particle i is denoted by \mathcal{C}_i . An unprotected particle has $\mathcal{P}_i \equiv \mathcal{C}_i$. In the case of hard spheres, $\mathcal{C}_i = \{\mathbf{r} \mid \|\mathbf{r} - \mathbf{r}_i\| \leq R_i\}$, the protective regions themselves are spheres of radius $R_i^{(\mathcal{P})} > R_i$ centered at $\mathbf{r}_i^{(\mathcal{P})}$. Thus, a spherical particle of radius R_i can be thought of as a *point particle* contained inside a protective sphere of radius $R_i^{(\mathcal{P})} - R_i$, i.e. $\|\mathbf{r}_i^{(\mathcal{P})} - \mathbf{r}_i\| \leq R_i^{(\mathcal{P})} - R_i$. In general it is not required that $\mathbf{r}_i^{(\mathcal{P})} = \mathbf{r}_i^0$, however, making the protective sphere concentric with the particle simplifies implementation.

The cube-shaped particles and protections described in detail in part I can be thought of as spheres but with an $L - 1$ distance metric function $\|\Delta\mathbf{r}\| = \max_{1 \leq k \leq d} \Delta r_k$, instead of the $L - 2$ Euclidean distance $\|\Delta\mathbf{r}\| = \sqrt{\sum_{k=1}^d \Delta r_k^2}$ characteristic of spheres. Except when the differences are essential, we will not explicitly distinguish between spheres and cubes with the understanding that distances are measured in an appropriate norm.

A. Mathematical Justification

The goal of the FPKMC algorithm is to generate a stochastic trajectory of the simulated system N Brownian particles with the correct probability. To achieve this goal the FPKMC algorithm samples the system's configuration $\mathbf{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_N\}$ at time t from the Green's function of the Master equation. In the considered case of particle diffusion, the Master

equation is a diffusion equation in the dN -dimensional space for the probability density $G(\mathbf{R}, t; \mathbf{R}_0)$ to find the system in configuration \mathbf{R} at time t , starting from \mathbf{R}_0 at time $t_0 = 0$,

$$\frac{\partial G}{\partial t} = \mathcal{D} \cdot \nabla^2 G \text{ in } \Omega, \text{ where } G(\mathbf{R}, 0) = \delta(\mathbf{R} - \mathbf{R}_0) \text{ and } G(\partial\Omega, t) = 0,$$

where $\mathcal{D} = \{D_1, \dots, D_N\}$ is a (block-diagonal) diffusion tensor. Here Ω denotes the region of phase space bounded by the non-overlap conditions between the hard particles. As soon as the configuration leaves region Ω through the bounding surface $\partial\Omega$ a hard-sphere collision occurs ¹.

What needs to be demonstrated is that the FPKMC algorithm correctly samples from $G(\mathbf{R}, t; \mathbf{R}_0)$ even if the latter Green's function is unknown. More specifically, the algorithm should correctly sample the time $\tilde{t} < t$ and configuration $\tilde{\mathbf{R}} \in \partial\Omega$ when the first collision occurs or, alternatively, sample a configuration \mathbf{R} given that no collision has occurred by time t .

Rather than trying to solve for the exact Green's function $G(\mathbf{R}, t; \mathbf{R}_0)$ - a difficult task indeed given the complicated shape of the collision boundary $\partial\Omega$ - the FPKMC algorithm relies on a much simpler Green's function obtained for an auxiliary sub-problem in which the particles are spatially protected against each other,

$$\frac{\partial G_0}{\partial t} = \mathcal{D} \cdot \nabla^2 G_0 \text{ in } \Omega_0 \text{ where } G_0(\mathbf{R}, 0) = \delta(\mathbf{R} - \mathbf{R}_0) \text{ and } G_0(\partial\Omega_0, t) = 0,$$

where now $\Omega_0 \subset \Omega$ is a Cartesian product of subspaces, each of which is a domain only of the coordinates of a single particle, $\Omega_0 = \mathcal{P}_1 \otimes \mathcal{P}_2 \otimes \dots \otimes \mathcal{P}_N$. This sub-problem is fully separable and can be decomposed into corresponding sub-problems for the Green's function $g_i(\mathbf{r}, t; \mathbf{r}_i^0)$ for each of the independent particles

$$\frac{\partial g_i}{\partial t} = D_i \nabla^2 g_i \text{ in } \mathcal{P}_i \text{ where } g_i(\mathbf{r}, 0) = \delta(\mathbf{r} - \mathbf{r}_i^0) \text{ and } g_i(\partial\mathcal{P}_i, t) = 0,$$

to give $G_0(\mathbf{R}, t; \mathbf{R}_0) = \prod_{i=1}^N g_i(\mathbf{r}_i, t; \mathbf{r}_i^0)$.

As elaborated in Ref. [16], using Green's theorem it can be shown that

$$G(\mathbf{R}, t; \mathbf{R}_0) = G_0(\mathbf{R}, t; \mathbf{R}_0) + \int_{\tilde{\mathbf{R}} \in \partial\Omega_0} \int_{\tilde{t}=0}^t \mathbf{J}_n(\tilde{\mathbf{R}}, \tilde{t}; \mathbf{R}_0) G(\mathbf{R}, t - \tilde{t}; \tilde{\mathbf{R}}) d\tilde{t} dS,$$

where $\mathbf{J}_n(\tilde{\mathbf{R}}, \tilde{t}; \mathbf{R}_0) = -\mathcal{D} \cdot \nabla_n G_0(\tilde{\mathbf{R}}, \tilde{t}; \mathbf{R}_0)$ denotes the normal component of the current at the auxiliary surface $\partial\Omega_0$. The above equation has an obvious probabilistic interpretation:

¹ Exactly how the particles react on collisions is outside the scope of the FPKMC algorithm.

the system's configuration at time t can be sampled by first sampling an exit time \tilde{t} and location $\tilde{\mathbf{R}} \in \partial\Omega_0$ on the auxiliary boundary. If $t < \tilde{t}$, the exit location is discarded and the system's configuration is sampled from (the normalized form of) $G_0(\mathbf{R}, t; \mathbf{R}_0)$. Otherwise one recursively repeats the process by constructing a new Ω_0 and samples an exit location and time starting from the new time origin, until the walk hits $\partial\Omega_0$ (this requires at least one pair protection). In this case, the exit location on the first-passage boundary $\partial\Omega_0$ serves as a source for subsequent propagation since the random walk is a memory-free Markov process that can continue from the last sampled point forgetting the past.

Owing to the separability of the protected walkers, the current is expressed as [16]

$$\mathbf{J}_n(\tilde{\mathbf{R}}, \tilde{t}; \mathbf{R}_0) = \sum_m \left(\prod_{l \neq m} g_l(\mathbf{r}_l, \tilde{t}; \mathbf{r}_l^0) \right) [-\nabla_n g_m(\tilde{\mathbf{r}}_m, \tilde{t}; \mathbf{r}_m^0)],$$

where now m denotes the walker that first leaves its protection through $\tilde{\mathbf{r}}_m \in \partial\mathcal{P}_m$ at time \tilde{t} . The probabilistic interpretation is as follows: first select a walker m and the associated $\tilde{t} = \tilde{t}_m$ with the correct probability, and then sample positions for all other walkers from (the normalized forms of) g_l 's. It is shown in Ref. [16] that m and the associated \tilde{t} can be sampled simply by choosing the minimum among independently sampled exit times, one time for each walker.

Numerical efficiency of the FPKMC algorithm, however, derives from an important additional observation not used in Ref. [16]: it is not necessary to sample g_l for the walkers that have not exited their protections. The exit times for all those walkers can be reused since their distributions have not changed. One can of course sample a new position $\mathbf{r}_l \in \mathcal{P}_l$ at time \tilde{t} , and then sample a new exit time from $g_l(\mathbf{r}, t - \tilde{t}; \mathbf{r}_l)$. However, as described in detail in part I, this will give an exit time and location with exactly the same distributions as those sampled originally. Therefore, those particles that have not exited their protections can be left in the queue or updated to the current point in time at will. For computational efficiency, the first-passage predictions for walkers are updated only when this becomes necessary, namely, for the walker m with the shortest exit time in the queue and, perhaps, for one or few other walkers whose protective regions need to be modified.

B. First-Passage Probability Densities

If the protective region \mathcal{P}_i of a particle i is disjoint from the protective regions of other particles, then the diffusive motion of the particle is independent of the motion of other particles, *as long as the particle is still inside its protection*. The motion of the particle inside its protection is a one-body diffusion problem that can often be solved analytically. The FPKMC algorithm entails sampling from the following two probability distribution functions (PDFs) for a particle initially at $\mathbf{r} = 0$ at time $t = 0$:

1. The *first-passage probability distribution* $J_1(\tilde{t}, \tilde{\mathbf{r}})$ that the particle first leaves its protective region at time \tilde{t} , when it is at position $\tilde{\mathbf{r}}$. We call the time \tilde{t} the *exit time* and the position $\tilde{\mathbf{r}}$ the *exit location*. For a point particle, $\tilde{\mathbf{r}}$ is on the surface of the protection, $\tilde{\mathbf{r}} \in \partial\mathcal{P}_i$.
2. The conditional *no-passage probability distribution* $c_1(\mathbf{r}; t)$ that the particle is at position $\mathbf{r} \in \mathcal{P}_i \setminus \partial\mathcal{P}_i$ at a given time t , *given* that it has not left its protection by time t .

These two probability distributions are the basic elements of the theory of first-passage processes [17] and are termed hereafter the first-passage and no-passage *propagators*, respectively. Appropriate solutions for the case of cube-shaped particles and protections were presented in the first part of this series of papers. Single-particle propagators for spherical particles inside spherical protection regions will be given in section III A.

1. Pair Propagators

In the FPKMC algorithm, the particles are protected by disjoint protective regions allowing the use of single-particle propagators to evolve the system. At some point in time, however, two particles i and j will collide and thus cannot be protected with disjoint regions. For efficient handling of collision events, nearly-colliding pairs are *associated* (partnered) and protected by a *pair protection region* \mathcal{P}_{ij} . We will focus on the case when $\mathcal{P}_{ij} = \mathcal{P}_i \cup \mathcal{P}_j$ consists of intersecting protections around each of the two particles, $\mathcal{P}_i \cap \mathcal{P}_j \neq \emptyset$. If either one of the particles leaves its protective region the pair *disassociates*.

Note that, for the case of additive hard spheres that we consider here, triple collisions never happen. Thus, it will always be possible to protect two colliding particles i and j as a pair even if there is a third particle k nearby. As the simulation time approaches the collision time, eventually i and j will be much closer to each other than to k and can thus be pair-protected with protective region \mathcal{P}_{ij} disjoint from \mathcal{P}_k . For non-additive hard spheres or other types of collision rules pair protection may not be sufficient and in such cases the approximate handling of interactions discussed in Section IV H will be useful.

For an associated pair of particles, in general, the FPKMC algorithm requires sampling from the following two distributions:

1. The first-passage probability distribution $J_2(\tilde{t}, \tilde{\mathbf{r}}_i, \tilde{\mathbf{r}}_j)$ that at time \tilde{t} , when the particles are at positions $\tilde{\mathbf{r}}_i$ and $\tilde{\mathbf{r}}_j$ respectively, one of two particles of the pair leaves its protection for the first time, or particles i and j collide and react for the first time.
2. The conditional no-passage probability distribution $c_2(\mathbf{r}_i, \mathbf{r}_j; t)$ for the positions of the particles at a given time t , *given* that neither i nor j has left its protection, nor a collision has occurred.

We term these two distributions *two-particle* or *pair propagators*. As discussed in part I of this series, the pair propagators allow further factorization into two single-body propagators. The first propagator is for the *difference walker* $\mathbf{r}_{ij}^{(D)} = \Delta\mathbf{r}_i - \Delta\mathbf{r}_j$, where the condition $\|\mathbf{r}_{ij}^{(D)}\| = R_i + R_j$ corresponds to a collision. The second propagator is for the *center walker* $\mathbf{r}_{ij}^{(C)} = w_i\Delta\mathbf{r}_i + w_j\Delta\mathbf{r}_j$, where w_i and w_j are appropriately-chosen weights. The protections and the propagators for the difference and center walkers are discussed in further detail in Section III B.

2. Hard-Wall Propagators

In some situations periodic boundary conditions may not be appropriate and instead hard-wall boundaries could be used along certain directions of the simulation box. Furthermore, additional absorbing or reflecting surfaces may be need to be inserted in the simulation volume, for example, to represent grain boundaries in a polycrystalline material. Particle and pair protections should be disjoint from any hard-wall surfaces or boundaries \mathcal{W} in order to ensure that single-particle and pair propagators can be used. However, some particle i will

eventually collide with \mathcal{W} and, thus, can not remain protected from \mathcal{W} forever. Therefore, particles near a wall \mathcal{W} are associated (paired) with the wall itself and protected with a *hard-wall protection* \mathcal{P}_i that is disjoint from all other single-particle and pair protections but intersects wall \mathcal{W} , $\mathcal{P}_i \cap \mathcal{W} = \partial\mathcal{P}_i^{\mathcal{W}} \neq \emptyset$. For such particles, we will use of the following *hard-wall propagators*:

1. The first-passage probability distribution $J_{HW}(\tilde{t}, \tilde{\mathbf{r}})$ that the particle exits its protective region or collides with the hard wall at time \tilde{t} , when it is at position $\tilde{\mathbf{r}}$.
2. The conditional no-passage probability distribution $c_{HW}(\mathbf{r}; t)$ that the particle is at position $\mathbf{r} \in \mathcal{P}_i \setminus \partial\mathcal{P}_i$ at a given time t , *given* that it has not left its protection or collided with the wall by time t .

C. Summary of FPKMC Algorithm

As already discussed in Section I B, the FPKMC algorithm processes a sequence of events in the increasing time order in an event loop. Each event has a *qualifier* (type of event) and is associated with a *primary* particle, which may have an event *partner* p_i . In the implementation of the algorithm described here the particles can only have one partner, either another particle forming a pair, or a hard wall. In general one can have multi-particle events and treat a group of particles together. We discuss this modification again in Section IV H and will report on its implementation details in future publications.

The most important types of events are first-passage propagations, in which a particle or pair reaches the boundary of its protective region or collides. Such events are scheduled and processed by sampling from appropriate first-passage distributions, and then processing any reactions, as appropriate. Following a propagation, a new protective region for the involved particle(s) is constructed that is as large as possible without overlapping with other existing protective regions. This may require destroying the protection of other (third party) particles or pairs whose protections block the construction of new protective regions. That is, if an unprotected particle gets too close to the protective region of another particle or pair, that particle or pair is brought to the current point in time using the appropriate no-passage propagator, and then new protective regions are constructed for all unprotected particles. One can keep a list of all unprotected particles or simply schedule immediate re-protection

events for them (i.e., place them at the head of the event queue).

Ideally, particles should be protected in such a way that the expected minimal first-passage time is maximized. That is, the protections have to be as large as possible, with more space given to particles with higher diffusion rates. If all particles are unprotected, the following method can be used to build optimal particle protections. First consider just two particles and protect them by two concentric non-overlapping protection domains Ω_1 and Ω_2 . The expectation time for the earliest of two exit times is maximal when domains Ω_1 and Ω_2 touch each other and their linear dimensions are proportional to the square roots of two diffusion coefficients $\sqrt{D_1}$ and $\sqrt{D_2}$. Obviously, in this case the expected first passage times for both particles are equal. Let us now define the distance between any two particles as the expectation for the their earliest exit time. First find the nearest neighbors in the whole system. If possible, protect them as a pair, otherwise, protect them individually with touching protective regions. Then, recursively protect the next nearest neighbors. A simple modification of this algorithm makes it applicable even when only a few particles are unprotected and one of them needs a new protection. One first finds the largest possible protection for that particle by performing a neighbor search over all nearby protected or unprotected particles. If the limiting neighbor is protected, one can simply protect the particle with a protection that touches the limiting protection. Otherwise, one recursively applies the same procedure to the unprotected limiting neighbor. Detailed pseudocode for this procedure is given in Section IV E.

However, it is highly non trivial to determine the optimal way to protect the particles in a dynamic context, when particles are created, destroyed, and move with widely differing diffusion coefficients. Specifically, following a single event, an optimal re-partitioning of the space into protective regions could be constructed, however, this would require destroying multiple existing protections and thus multiple propagations and new event predictions. Obviously, it is necessary to find a reasonable trade-off between re-building too many protections and giving all of the particles sufficient room to move over sufficient distances compared to other particles. Intuitively, the algorithm's performance will be optimized by giving more protective space to the particle or a small subset of particles that dominate(s) the event queue, up to the point when giving more room to these *fast particles* no longer increases the average time interval between subsequent events (or actually makes efficiency worse by squeezing the other (slow) particles too much). In Section IV G we will describe space partitioning strate-

gies we applied in several specific situations encountered in FPKMC simulations. While some of our strategies are relatively general and have tuning parameters to cater to different specific cases, understanding of the problem at hand helps considerably in optimizing the algorithm performance.

We conclude this section with a few additional notes on the components of the FPKMC algorithm. Whenever an event creates new particles, a check is made to determine whether the newly inserted particle overlaps with any existing particles. If it does, appropriate reactions are immediately processed and the overlap check is repeated until no overlap is detected. Additionally, unlikely the mobile particles, the immobile particles remain unprotected and have no partners, but other particles can have them as partners. This is particularly useful when there are large immobile particle (e.g. clusters of monomers) that are surrounded by a dense pool of small mobile particles (e.g., monomers). The mobile particles do not affect each other for as long as their protections do not overlap, even if they share the same immobile cluster as a partner. However, when an event changes the configuration of the immobile cluster (e.g., it absorbs a monomer), all its partner particles are brought to the current time and re-protected.

III. SINGLE PARTICLE AND PAIR PROPAGATORS

In this section we describe first-passage and no-passage propagators for spherical particles. The propagators for single particles are similar to the ones described in paper I of this series for cube-shaped particles and are discussed here only briefly. The reader is referred to Appendices A and B of paper I for more information. However, the pair propagators for two spherical particles are considerably different than those for cubical particles and are presented in more detail.

A. Single Particle Propagators

The FPKMC algorithm requires first-passage and no-passage propagators for a single spherical particle A of radius R_A with diffusion coefficient D_A , starting from an initial position \mathbf{r}_A^0 at the center of a protective sphere \mathcal{P}_A of radius $R_A^{\mathcal{P}} > R_A$ concentric with the particle. This reduces to finding the propagators for a *point* Brownian particle starting at

time $t = 0$ from the center of sphere with radius $R_P^{\mathcal{P}} = R_A^{\mathcal{P}} - R_A$. Due to the full rotational symmetry, the first-passage PDF $J_1(\tilde{t})$ is a function of time only, and the exit location on the protective sphere \mathcal{P}_A can be sampled from a uniform distribution. Similarly, the no-passage PDF $c_1(r; t)$ becomes a function of time and radial distance only, while the exit position can be sampled from a uniform distribution on the sphere of radius r .

Expressing the distances and time in the reduced units of R and R^2/D , the propagators can be obtained by solving the diffusion problem for a point Brownian particle with diffusion coefficient $D = 1$ inside a sphere of unit radius $R_P^{\mathcal{P}} = 1$,

$$\frac{\partial [rc(r, t)]}{\partial t} = \frac{\partial^2 [rc(r, t)]}{\partial r^2}$$

with the boundary condition $c(1, t) = 0$ and the initial condition $c(r, 0) = \delta(r)$. We focus on two useful series expansions of the solution, one that converges quickly at short times ($t \lesssim 1/4$),

$$c(r, t) = (4\pi t)^{-\frac{3}{2}} \sum_{m=-\infty}^{\infty} \left(1 + \frac{2m}{r}\right) \exp\left[-\frac{(r + 2m)^2}{4t}\right] \quad (2)$$

and another that converges quickly at long times ($t \gtrsim 1/\pi^2$),

$$c(r, t) = \frac{1}{2r} \sum_{m=1}^{\infty} m \sin(m\pi r) e^{-m^2\pi^2 t}. \quad (3)$$

Below we use these two solutions to obtain the propagators and to construct over-estimators and bounds necessary for efficient rejection sampling.

1. First-passage propagator

Integration of $c(r, t)$ over the unit sphere yields the survival probability in two infinite series forms, a short time series

$$S(t) = S_{n \rightarrow \infty}(t), \quad (4)$$

where

$$S_n = (2\pi t^3)^{-\frac{1}{2}} \sum_{m=-n-1}^n \left[2m + \frac{(1 + 2m)^2}{2t}\right] \exp\left[-\frac{(1 + 2m)^2}{4t}\right],$$

and a long time series

$$S(t) = -2\pi^2 \sum_{m=1}^{\infty} (-1)^m m^2 e^{-m^2 \pi^2 t}. \quad (5)$$

In order to use a rejection sampling procedure similar to the one described in Appendix A of paper I, we select a switchover time τ such that $1/\pi^2 \leq \tau \leq 1/4$ and use the piecewise smooth function $C(t)$ as an over-estimator for the survival probability at all times,

$$C(t) = \begin{cases} \frac{S(\tau)}{S_s(\tau)} S_s(t) & \text{for } t < \tau \\ S_l(t) & \text{for } t > \tau, \end{cases}$$

where $S_s(t) = (\frac{1}{t} - 2)e^{-\frac{1}{4t}} / \left[(2\pi)^{\frac{1}{2}} t^{\frac{3}{2}} \right]$ and $S_l(t) = 2\pi^2 e^{-\pi^2 t}$ are the leading terms of the short- and long-time series for the survival probability. A sample exit time t is obtained by solving $C(t) = \xi$, where ξ is random number uniformly distributed in $[0, 1)$. Solving this equation for $t < \tau$ can be done efficiently by Newton iteration. To obtain the sequence of converging bounds necessary for rejection sampling, let us note that the terms of the long-time series alternate in sign and decrease in magnitude with increasing m and can themselves serve as the needed bounds. For the short-time series we can bound the remainder $R_n(t) = S(t) - S_n(t)$ with

$$0 \leq R_n(t) \leq (4\pi t^3)^{-\frac{1}{2}} \left[\frac{(1+2m)^2}{t} - 1 + 2m \right] \exp \left[-\frac{(1+2m)^2}{4t} \right].$$

In our implementation we use $\tau = 0.243$, which gives a rejection ratio of about 0.6%. The maximum relative error in the survival probability is $C(t)/S(t) - 1 \approx 7 \cdot 10^{-3}$, so that $C(t)$ is an acceptable approximation for $S(t)$ without rejection sampling.

2. No-passage propagator

To enable efficient rejection sampling, we need a tight over-estimator $C(r, t)$ for the no-passage PDF at arbitrary time t . Here we construct such a function by stitching together two different expressions appropriate for times shorter and longer than a switchover time τ ($1/\pi^2 \leq \tau \leq 1/4$). For times $t < \tau$ the leading term ($m = 0$) in the short-time series solution (1) is a reasonable over-estimator:

$$C_s(r, t) = (4\pi t)^{-\frac{3}{2}} e^{-\frac{r^2}{4t}}.$$

On the other hand, for times $t > \tau$ a good over-estimator is given by

$$C_l(r, t) = \frac{\sin \pi r}{2r} e^{-\pi^2 t} \left[1 + \frac{1}{2\pi^2 t} + \frac{1}{2\sqrt{\pi t}} \operatorname{erfc}(\pi^2 t) \right].$$

Sampling from $C_s(r, t)$ entails evaluation of one inverse error function while sampling from $C_l(r, t)$ requires solving for r the following equation

$$\frac{1}{\pi} \sin \pi r - r \cos \pi r = \xi,$$

where ξ is a random number uniformly distributed on $[0, 1)$. The solution can be found by Newton iteration using $r = \sqrt[3]{\frac{3\xi}{\pi^2}}$ as a starting guess if $\xi \leq 1/2$ and $r = 1 - \frac{1}{\pi} \sqrt{2(1 - \xi)}$ if $\xi > 1/2$.

The same two series solutions can be used to obtain a sequence of increasingly tight bounds on the no-passage PDF. For short times, we approximate $c(r, t)$ with the partial sum

$$c_m(r, t) = \sum_{k=-m}^m \left(1 + \frac{2k}{r} \right) \exp \left[-\frac{(r + 2m)^2}{4t} \right],$$

and use the following lower and upper bounds for the remainder term, $R_l \leq c(r, t) - c_m(r, t) \leq R_u$,

$$\begin{aligned} R_l &= \left(1 - 2\frac{m+1+t}{r} \right) \exp \left[-\frac{(r - 2m - 2)^2}{4t} \right] + \frac{2t}{r} \exp \left[-\frac{(r + 2m + 2)^2}{4t} \right] \\ R_u &= \left(1 + 2\frac{m+1+t}{r} \right) \exp \left[-\frac{(r + 2m + 2)^2}{4t} \right] + \frac{2t}{r} \exp \left[-\frac{(r - 2m - 2)^2}{4t} \right], \end{aligned}$$

For long times, efficient bounds on the no-passage PDF are obtained from the following bound for the magnitude of the remainder

$$|4\pi r^2 \sum_{k=m+1}^{\infty} \frac{k}{2r} \sin(k\pi r) e^{-m^2 \pi^2 t}| \leq \left[2\pi r(m+1) + \frac{r}{\pi t} \right] e^{-(m+1)^2 \pi^2 t}.$$

B. Pair Propagators

In this section we discuss particle protection and propagators used to enable collisions of pairs of particles. Consider two spherical Brownian particles A and B of radii R_A and R_B and diffusion coefficients D_A and D_B . Each of the particles is protected by a sphere of radii $R_{A/B}^P > R_{A/B}$ concentric with the initial particle positions $\mathbf{r}_{A/B}^0$. When two protective

spheres overlap, i.e., $R_A^{\mathcal{P}} + R_B^{\mathcal{P}} > r_{AB} = \|\mathbf{r}_A - \mathbf{r}_B\|$, the particles can collide while diffusing within their individual protections. Sampling of particle collisions is enabled by transforming the two-particle diffusion problem into two single-particle problems, one for the *difference walker* $\mathbf{r}_D = \mathbf{r}_A - \mathbf{r}_B$ and the other for the *center walker* $\mathbf{r}_C = w_A \mathbf{r}_A + w_B \mathbf{r}_B$. A collision occurs when the difference walker reaches the collision radius $r_D = \|\mathbf{r}_A - \mathbf{r}_B\| = R_A + R_B$. It can be shown that, with the choice $w_A D_A = w_B D_B$, the cross term $\frac{\partial^2}{\partial \mathbf{r}_C \partial \mathbf{r}_D}$ in the Laplacian operator vanishes which means that the six-dimensional PDF for the pair factorizes into a product of two three-dimensional PDFs, one for each walker. The resulting center and difference walkers diffuse independently with diffusion coefficients $D_D = D_A + D_B$ and $D_C = w_A^2 D_A + w_B^2 D_B$.

There is considerable freedom for choosing weights w_A and w_B and protections for the center and difference walkers. To simplify the implementation, the center and difference walkers are each protected by spheres of radii $R_D^{\mathcal{P}}$ and $R_C^{\mathcal{P}}$ centered around their initial positions. It turns out that choosing $R_D^{\mathcal{P}} = R_C^{\mathcal{P}}$ maximizes the use of space within the overlapping protective spheres of the two original particles. Furthermore, by setting $D_C = D_D$ the expectation values for the exit times for the difference and center walkers become similar and nearly maximal. This leads to the following optimal choice for the coordinate transformation

$$\begin{aligned} \mathbf{r}_D &= \mathbf{r}_A - \mathbf{r}_B \\ \mathbf{r}_C &= \sqrt{\frac{D_B}{D_A}} \mathbf{r}_A + \sqrt{\frac{D_A}{D_B}} \mathbf{r}_B. \end{aligned}$$

Finally, the condition that a collision should be possible requires that $R_C^{\mathcal{P}} > \delta$, where $\delta = r_{AB} - (R_A^{\mathcal{P}} + R_B^{\mathcal{P}})$ is the initial inter-particle gap. We take

$$R_C^{\mathcal{P}} = R_D^{\mathcal{P}} = (2 + \alpha)\delta,$$

where $\alpha \geq -1$ is a parameter in the algorithm. We find that $\alpha = 1$ or $R_C^{\mathcal{P}} = 2\delta$ is a reasonable choice but, when there is more room available, one can increase α up to a maximal value α_{max} specified in the code. It can be shown that, as the difference and center walkers propagate inside their own protections, each of the two original particles remains within a sphere of radius

$$R_{A/B}^{\mathcal{P}} = R_{A/B} + (2 + \alpha)\delta \sqrt{D_{A/B}} \frac{\sqrt{D_A} + \sqrt{D_B}}{D_A + D_B}.$$

This defines the size of the protective region around each original particle that is minimally necessary to allow pair protection, i.e., the minimal distance to the next-nearest neighbors of particles A and B which allows for pair protection.

1. The Difference and Center Propagators

With the above coordinate transformation, the first-passage problem for the pair of original particles separates into two independent first-passage problems for the center walker and the difference walker. The overall first-passage time \tilde{t} is the smaller of the two first-passage times. Thus, the possible first-passage events for the pair are:

Collision when the difference walker reaches the surface $r_D = \|\mathbf{r}_A - \mathbf{r}_B\| = R_A + R_B$.

Dissolution when the difference walker reaches the surface $r_D = R_D^{\mathcal{P}}$. The center walker has not yet left \mathcal{P}_C and can be updated using the single-point no-passage propagator $c_1(\Delta\mathbf{r}_C; \tilde{t})$.

Displacement when the center walker reaches the surface $r_C = R_C^{\mathcal{P}}$. The difference walker has not yet left \mathcal{P}_D and can be updated using a special single-point no-passage propagator $c_D(\Delta\mathbf{r}_D; \tilde{t})$ (see the next section).

Because the center walker is restricted inside a spherical protection \mathcal{P}_C of radius $R_C^{\mathcal{P}}$, the same FP and NP propagators described earlier in Section III A for single particles can be used for this walker.

Unlike the case of cube-shaped protections considered in paper I, protection region of the difference walker for a pair of spherical particles cannot be a sphere. Among several choices considered, in our current implementation we opt to protect the difference walker with a *cut sphere*, i.e. $\mathcal{P}_D = \{\mathbf{r}_D \mid \|\Delta\mathbf{r}_D\| < R_D^{\mathcal{P}} \text{ and } \|\mathbf{r}_D\| \geq R_A + R_B\}$, where $\Delta\mathbf{r}_D = \mathbf{r}_D - \mathbf{r}_D^0$. Lacking the spherical symmetry and having more complex geometry of the first-passage surface (see Fig. 1), finding a usable analytical solution for the propagators in this protective volume is more involved than for the cubes. Various approximations for the distribution of exit times and locations $J_D(\tilde{t}, \tilde{\mathbf{r}}_D)$ for cut spheres have been considered in the context of diffusion Monte Carlo [18, 19]. However, the previous work used only time-averaged solutions but not the full time-dependent Green's function $c_D(\mathbf{r}_D; t)$. In the next section we




Figure 1: Protective region \mathcal{P}_D for the difference walker in the case of two hard disks in two dimensions. In this case \mathcal{P}_D is a cut disk (blue) and the propagator $c_D(\mathbf{r}_D; t)$ is analytically complex.

describe a practical alternative to the full analytical solution of the first-passage problem inside the cut sphere.

2. Hopping-Based Propagators

Rather than trying to solve for the exact propagators for the center walker in a cut sphere, we resort to generating a random walk through a sequence of small displacements in randomly chosen directions. In principle, restricting the hops to sufficiently small displacements, such discrete walks can approximate the samples from the exact continuum distributions to any desired accuracy. In practice, we use random walks only for particle pairs that have been brought close to collisions through the use of exact continuum single-particle propagators. For such pairs, the walks are typically short and do not entail high computational cost.

The time when the trajectory brings the difference walker to the boundary $\partial\mathcal{P}_D$ of the cut sphere \mathcal{P}_D is taken as an approximation for the first-passage time \tilde{t} and the location where the sampled walk hits the surface $\partial\mathcal{P}_D$ is an approximation to $\tilde{\mathbf{r}}_D$. At every hop a constant increment $\Delta t_h = \Delta r_h^2 / (2D_D)$ is added to the running time and the displacements along each dimension are sampled independently from a one-dimensional Gaussian probability distribution with the standard deviation Δr_h . Alternatively, the magnitudes of walker displacements can be kept equal to Δr_h while the time increments can be sampled from the first-passage distribution J_1 for the sphere of radius Δr_h . In the limit of small displacement

length both methods should reproduce the required first-passage probability distributions in the cut sphere.

As a practical matter, one must ensure that the path makes some progress, that is at least one hop is taken. Therefore Δr_h should be smaller than some fraction (e.g., half or one third) of the inter-particle gap δ . Furthermore, to ensure reasonable accuracy the displacement Δr_h should be smaller than the relevant length-scales of the cut sphere \mathcal{P}_D . We set $\Delta r_h = \epsilon_h \min(\delta, R_A + R_B)$, where δ is the initial inter-particle gap and $\epsilon_h \ll 1$ is a fractional *hop length* parameter. The walk terminates when a collision or pair dissolution occurs, or when time exceeds a specified upper bound on \tilde{t} , t_{max} . Occasionally, the first displacement has to be truncated in order to ensure that the difference walker does not leave \mathcal{P}_D after only one hop: we simply re-sample the hop again if the initial displacement is too large. In cases when the particles reflect on collisions, we simply reject the last hop leading to collision. In cases when the particles react on collisions, we take the full length of the last hop but truncate the last time increment. If instead of collision a pair dissolution event occurs, we reject the last hop rather than trying to truncate it more accurately (and expensively). In all cases, exit time $\tilde{t}_D = t$ and location $\tilde{\mathbf{r}}_D$ are recorded for possible future use. If the time t exceeds t_{max} , the last time increment is set to $\Delta t'_h = t - t_{max}$ and the length of the last hop is scaled by factor $\sqrt{\Delta t'_h / \Delta t_h}$.

If and when the previously scheduled first-passage event actually occurs, we simply move the difference walker to the pre-sampled exit location $\mathbf{r}_D = \tilde{\mathbf{r}}_D$. However, when the scheduled first-passage event is preempted by some other event, e.g. intrusion of a third party particle, a no-passage propagation is required instead. In such instances, rather than storing the entire walk generated during pre-sampling, we opt to repeat the same trajectory starting from the same seed for the pseudo-random number generator. For this purpose we store the random seed for each protected pair when its next first-passage event is scheduled, and use this seed should it become necessary to repeat the previously sampled walk.

IV. ALGORITHMIC DETAILS

In this section we give details of various components of the FPKMC algorithm other than the propagators. First, we briefly discuss computational techniques for searching for near neighbors. We then explain the types of events that are scheduled and processed by the

FPKMC algorithm. This section contains several detailed pseudocodes, including the main event loop and the construction of protective regions.

A. Near-Neighbor Search

Efficient particle-based algorithms use various geometric techniques to reduce to $O(1)$ the cost of searching for the neighbors of a given particle. Reference [20] provides extensive details and illustrations of these techniques for hard spheres and ellipsoids; here we summarize only the essential components.

1. *Linked List Cell (LLC) Method*

The most basic technique is the so-called *linked list cell* (LLC) method. The simulation domain, typically an orthogonal box, is partitioned into N_c cells, typically cubes. Each particle i stores the cell c_i to which its centroid belongs, and each cell c stores a list \mathcal{L}_c of all the particles it contains. Given a particle and a search range, the lists of potential neighbors is determined by scanning through the neighboring cells. Typically, for maximal efficiency the cell should be larger than the largest search range, that is, larger than the largest protective region \mathcal{P}_i . If some particle grows too large, e.g. due to coalescence reactions, one can enlarge the cells and re-build the associated linked lists. In Section IV G we discuss the impact of cell size on efficiency as well as methods for choosing the optimal cell size in simulations of radiation damage.

In our FPKMC implementation, in addition to the list of particles \mathcal{L}_c , each cell c stores a *bitmask* \mathcal{M}_c consisting of $N_{\text{bits}} > N_s$ bits (N_s is the number of species). Bit γ in the bitmask \mathcal{M}_c is set if cell c contains a particle of species γ . The bit is set whenever a particle of species γ is added to the cell, and all bitmasks are cleared periodically. When performing a neighbor search for particle i , cells not containing particles of species that interact with species s_i are easily found and are simply skipped. This can significantly speed up the neighbor searches in cases where not all particles interact with all other particles, for example in simulations of two species diffusion-limited annihilation $A + B \rightarrow 0$ where particles of like species do not interact.

2. Near-Neighbor List (NNL) Method

Another neighbor search technique is the *near-neighbor list* (NNL) method [20]. The essential idea is to enclose every protective region \mathcal{P}_i inside a *bounding neighborhood* \mathcal{N}_i , $\mathcal{P}_i \subset \mathcal{N}_i$. This bounding neighborhood is fixed while the particle and its protection change as the particle moves around, until the particle comes close to the boundary of \mathcal{N}_i at which point a new \mathcal{N}_i is constructed. In principle, one can treat the boundary of \mathcal{N}_i as a first-passage surface, however we simply rebuild the bounding neighborhood whenever the particle comes close to its boundary.

The linked list $\text{NNL}(i)$ lists all other neighborhoods that intersect neighborhood \mathcal{N}_i (hard walls or other boundaries may also be near neighbors). This is used to identify potential *interactions* of particle i and can be reused until the particle core \mathcal{C}_i approaches the boundary of \mathcal{N}_i . This results in significant savings of computational effort if particle motion is localized and the particles experience numerous displacements before leaving their bounding neighborhood. Note that the LLC method is still used to build \mathcal{N}_i and $\text{NNL}(i)$ which keeps the maximal cost of pairwise searches at $O(N)$ instead of $O(N^2)$. In our implementation \mathcal{N}_i is build as a sphere concentric with the particle and with the diameter larger than $2\mu R_i$ but smaller than the cell size, where $\mu > 1$ is a parameter.

Even the NNL method can become inefficient when some particles are much larger than others, e.g. large clusters formed by coalescence of defects in radiation damage modeling. In such cases, if the cells are still maintained larger than the size of the largest particle, the same cells may contain many small particles making the search for near neighbors expensive. On the other hand, if the cells are kept reasonably small, neighbor searches need to examine many small cells in order to account for all near neighbors of the large particles. The idea of *bounding sphere complexes* (BSCs) [20] method is to use small cells but cover \mathcal{N}_i with a collection of N_{BSC} spheres not larger than the cells. The so-constructed sphere complex remains immobile until \mathcal{N}_i changes, which would occur infrequently if the large particles move slowly or are immobile. The small spheres forming BSCs are inserted in the LLCs and near neighbors of each large particle are found by searching for the near neighbors of each constituent small sphere in the corresponding BCS.

B. Types of Events

The following types of events, as identified by the event partner p and event type ν (which we arbitrarily represent with an integer here), are scheduled and processed in the FPKMC algorithm for a given particle i of species α :

Particle protection identified by $p = 0, \nu = 1$. Such events are scheduled at the beginning for all particles and immediate updates are scheduled for any particles whose protections are destroyed before the scheduled first-passage time. Processing consists of protecting the particle with a new protective region \mathcal{P}_i (see Section IV E).

Particle update identified by $p = 0, \nu = 0$. A protected particle needs a new event prediction.

Particle insertion identified by $p = 0, \nu = -1$. Processing requires first checking whether the newly inserted particle overlaps with any existing particles. If it does, corresponding reactions are processed, otherwise, the particle is protected (see Section IV D below).

First-passage hop identified by $p = i, \nu < 0$. Scheduling consists of sampling an exit time \tilde{t} and exit location $\tilde{\mathbf{r}}$ from J_1 . Processing consists of moving the particle, $t_i \leftarrow t_i + \tilde{t}$, $\mathbf{r}_i \leftarrow \mathbf{r}_i + \tilde{\mathbf{r}}$, destroying the old protection, and then protecting the particle again. Additional information, such as the surface of \mathcal{P}_i with which the particle collides, can be recorded in ν .

Particle decay identified by $p = i, \nu > 0$. The particle decays via the reaction \mathcal{D}_ν^α before leaving its protection. Scheduling consists of sampling an exponentially-distributed number with mean $(\sum_k \Gamma_k^\alpha)^{-1}$ and then choosing one of the reactions ν with probability $\Gamma_\nu^\alpha / \sum_k \Gamma_k^\alpha$ (here the sum is over all single particle decay reactions available for particle i). Processing consists of sampling from $c_1(\mathbf{r}; t)$ and moving the particle, $t_i \leftarrow t$, $\mathbf{r}_i \leftarrow \mathbf{r}_i + \mathbf{r}$, destroying the protection \mathcal{P}_i , and then executing the selected decay reaction \mathcal{D}_ν^α .

Hard-wall update identified by $p = -w, \nu = 0$, where w is the identifier of the hard wall. A particle-wall pair needs a new event prediction after it has been protected.

Hard-wall escape identified by $p = -w$, $\nu < -1$. This event is similar to a first-passage hop, however, J_{HW} is sampled instead of J_1 . Additional information such as the surface of \mathcal{P}_i with which the particle collides could be recorded in ν .

Hard-wall collision identified by $p = -w$, $\nu = -1$. This event is similar to the hard-wall escape, however, the outcome is a collision of the particle with the hard wall resulting in annihilation of the particle (a reflection is never explicitly scheduled).

Hard-wall decay identified by $p = -w$, $\nu > 0$. The particle decays via the reaction \mathcal{D}_ν^α before colliding with its protection or the hard wall.

Pair update identified by $p = j$, $\nu = 0$. A pair needs a new event prediction after it has been protected.

Pair disassociation identified by $p = j$, $\nu < -1$, where j is the partner particle. Scheduling consists of sampling from $J_2(\tilde{t}, \tilde{\mathbf{r}}_i, \tilde{\mathbf{r}}_j)$. Processing consists of moving both particles $t_i \leftarrow t_i + \tilde{t}$, $\mathbf{r}_i \leftarrow \mathbf{r}_i + \tilde{\mathbf{r}}_i$, and $t_j \leftarrow t_j + \tilde{t}$, $\mathbf{r}_j \leftarrow \mathbf{r}_j + \tilde{\mathbf{r}}_j$, destroying both protections \mathcal{P}_i and \mathcal{P}_j , scheduling an immediate protection event for particle j , and then protecting particle i again. The specific meaning of the disassociation event could be recorded in ν , e.g. that particle j left its protection or that the center walker left its protection.

Pair collision identified by $p = j$, $\nu = -1$. Scheduling consists of sampling from $J_2(\tilde{t}, \tilde{\mathbf{r}}_i, \tilde{\mathbf{r}}_j)$. Processing consists of executing reaction $\mathcal{R}_{\alpha\beta}$, where β is the species of partner particle j .

Pair decay identified by $p = j$, $\nu > 0$. Particle i decays before the pair disassociates or collides. Processing consists of sampling $c_2(\mathbf{r}_i, \mathbf{r}_j; t)$ and moving both particles accordingly, destroying both protections, scheduling an immediate protection for particle j , and then processing the decay reaction \mathcal{D}_ν^α for particle i .

C. Main Event Loop

The core of the FPKMC algorithm is the event loop described in Algorithm 2. We have already discussed most of its steps and give here additional details to facilitate other implementations of the FPKMC algorithm:

- Immobile particles may have multiple partners, which are not stored explicitly. Searching for partners of an immobile particle i consists of performing a neighbor search for any protective regions \mathcal{P}_j that might overlap with \mathcal{C}_i and checking if $p_j = i$.
- For pairs ij , only one of the particles is inserted into the heap, specifically, $\min(i, j)$ if both are mobile or just the mobile particle otherwise. In our implementation deletions from the event queue heap are lazy, meaning that particles that should be deleted from the heap [for example, $\max(i, j)$ for a pair of mobile particles] are still in the heap but with infinite timestamp.
- Un-protecting a particle i consists of bringing it and, potentially, its partner to the current time using propagators c_1 , c_2 or c_{HW} depending on partner type p_i , and scheduling immediate protection events for i and its partner.
- Whenever the position of a particle is updated, its time is updated as well, $t_i \leftarrow t$, and also the LLCs and NNs are updated accordingly.

Algorithm 2: FPKMC event loop. Here ξ denotes a uniform random variate $0 < \xi < 1$. Initially the time of the next particle insertion $t_B = -1$, and all particles are put in the event queue with $t_e = 0$, $p = 0$, $\nu = 1$.

1. Find (query) the top of the event heap to find the next particle i to have an event with p at t_e .
2. If $t_B < 0$, set $t_B = -\ln \xi / \Gamma_B$, where $\Gamma_B = \sum_{\alpha=1}^{N_s} \mathcal{B}_\alpha$ is the total insertion rate of all particle types.
3. If $t_B < t_e$, set $t \leftarrow t_B$, $t_B \leftarrow -1$, insert a new particle into the system (Algorithm 3) and cycle back to step 1.
4. Remove particle i from the event queue, store the time increment $\Delta t = t_i^e - t_i$, the event partner $p = p_i$, event type $\nu = \nu_i$, and particle species $\alpha = s_i$. Advance the simulation time $t \leftarrow t_i^e$.
5. *Insertion:* If $p = 0$, $\nu = -1$, check if the newly inserted particle overlaps with any existing particle and if so, process corresponding reactions. If the particle still exists after the check, set its event to be a regular protection, $p \leftarrow 0$, $\nu \leftarrow 1$. Otherwise, cycle back to step 1.

6. *Single-particle event*: If $p = i$ and $\nu \neq 0$, update particle i :
 - (a) *First passage hop*: If $\nu < 0$, propagate i to its first-passage location, $\mathbf{r}_i \leftarrow \mathbf{r}_i + \tilde{\mathbf{r}}$, and destroy its protection \mathcal{P}_i . For spherical particles, one can sample $\tilde{\mathbf{r}}$ at this point instead of pre-computing it when scheduling first-passage hops (see Section III A).
 - (b) *Decay*: Else if $\nu > 0$, sample \mathbf{r} from $c(\mathbf{r}; \Delta t)$, set $\mathbf{r}_i \leftarrow \mathbf{r}_i + \mathbf{r}$, destroy \mathcal{P}_i , process the decay reaction \mathcal{D}_ν^α and cycle back to step 1.
7. *Hard-wall event*: Else if $p = -w < 0$ and $\nu \neq 0$, update the particle-wall pair $p - w$:
 - (a) *First-passage event*: If $\nu < 0$, propagate particle i to its first-passage time, update \mathbf{r}_i and ν_i if needed. *Collision*: If $\nu = \nu_i = -1$, delete particle i and cycle back to step 1.
 - (b) *Decay*: Else if $\nu > 0$, sample \mathbf{r} from $c_{HW}(\mathbf{r}; \Delta t)$, set $\mathbf{r}_i \leftarrow \mathbf{r}_i + \mathbf{r}$, destroy \mathcal{P}_i , process the decay reaction \mathcal{D}_ν^α and cycle back to step 1.
8. *Pair event*: Else if $p > 0$ and $p \neq i$ and $\nu \neq 0$, update the particle pair ij , $j = p$, $\beta = s_j$. Test whether the partner is a mobile particle, i.e., whether $p_j = i$.
 - (a) *Disassociation*: If $\nu < 0$, propagate the pair to its first passage time (see Section III B), update \mathbf{r}_i and, if needed, \mathbf{r}_j and ν_i . Destroy protections \mathcal{P}_i and \mathcal{P}_j .
 - (b) *Collision*: If $\nu = \nu_i = -1$, process the particle-particle collision. If j is immobile, find and un-protect all of its other partners k (see Section IV F). Process reaction $\mathcal{R}_{\alpha\beta}$, scheduling immediate particle insertion events for any remaining or newly created particles, and cycle back to step 1.
 - (c) *Disassociation*: Else if $\nu < 0$, schedule an immediate protection for the partner, $p_j \leftarrow 0$, $\nu_j \leftarrow 1$, and insert j into the event queue.
 - (d) *Decay*: Else if $\nu > 0$, sample $c_2(\mathbf{r}_A, \mathbf{r}_B; \Delta t)$, set $\mathbf{r}_i \leftarrow \mathbf{r}_i + \mathbf{r}_A$ and destroy \mathcal{P}_i . If j is mobile, set $\mathbf{r}_j \leftarrow \mathbf{r}_j + \mathbf{r}_B$, destroy \mathcal{P}_j , schedule an immediate protection for the partner, $p_j \leftarrow 0$, $\nu_j \leftarrow 1$, and insert j into the event queue. Process the decay reaction \mathcal{D}_ν^α and cycle back to step 1.
9. If particle i is mobile, build a new protection \mathcal{P}_i and find the limiting neighbor k (Algorithm 4):

- (a) If $\nu \neq 0$ and $k \equiv \mathcal{N}_i$, then rebuild \mathcal{N}_i and $\text{NNL}(i)$.
 - (b) Else if $\nu \neq 0$ and k is a protected particle, then try to enlarge \mathcal{P}_i by making more room for it:
 - i. Destroy \mathcal{P}_i and un-protect k .
 - ii. Build a new protection \mathcal{P}_i and find the new limiting neighbor k' .
 - iii. If k' is a protected particle, un-protect k again, set $k \leftarrow k'$ and cycle back to step 9(b)i.
10. Let $p = p_i$ (this may have changed during the previous step). If $p = i$ or $p = 0$ then schedule a new single-particle event for i :
- (a) If $N_d^\alpha > 0$ then sample the time of next decay $t_d = -\ln r / \Gamma_d^\alpha$, where $\Gamma_d^\alpha = \sum_k \Gamma_k^\alpha$, and find reaction ν for which $\sum_{k=1}^{\nu-1} \Gamma_k^\alpha < r \Gamma_d^\alpha < \sum_{k=1}^\nu \Gamma_k^\alpha$ using linear or binary search. Otherwise let $t_d = \infty$.
 - (b) If i is protected, sample an exit time \tilde{t} and (optionally) an exit location $\tilde{\mathbf{r}}$ from J_1 using t_d as an upper bound. Otherwise set $\tilde{t} = \infty$.
 - (c) Choose the smaller of t_d and \tilde{t} and insert i in the event queue with the appropriate event prediction.
11. Else if $p > 0$ and $p \neq i$, schedule a new pair event for ij , $j = p$, $\beta = s_j$. If j is mobile, i.e., if $p_j = i$, then set $k = \min(i, j)$, otherwise set $k = i$.
- (a) Sample a new decay time for particle i , t_i^d and sample t_j^d if j is mobile. Set t_k^e to the time of the first decay reaction, set ν_k to the selected decay reaction and set $k_d = i$ or $k_d = j$ to indicate the decaying particle.
 - (b) Sample an exit time and location for the pair from $J_2(\tilde{t}, \tilde{\mathbf{r}}_i, \tilde{\mathbf{r}}_j)$ using t_k^e as an upper bound. If $\tilde{t} < t_k^e$, change $t_k^e \leftarrow \tilde{t}$ and change ν_k appropriately. If needed, store the exit location.
 - (c) If $\nu_k > 0$, insert k_d into the event queue with the appropriate decay event prediction. If $k_d \neq k$, delete k from the event queue.

12. Else if $p = -w < 0$ then schedule a new event for the particle-wall pair similarly to the case of an immobile partner in step 11.
 13. Cycle back to step 1.
-

D. Particle Insertion

Pseudocode 3 is an algorithm for inserting a new particle in the simulation volume. Its inverse operation of particle deletion is not given here for brevity. Note that in some applications it can be necessary to insert a whole collection of particles rather than a single particle. For example, irradiation by heavy particles - ions or neutrons - creates large displacement cascades that quickly anneal into a whole collection of point defects and clusters. The steps for inserting particles groups are similar to those for inserting a single particle and are not described here.

Algorithm 3: Inserting a newly-born particle.

1. Choose which particle species α is to be inserted with probability $\mathcal{B}_\alpha/\Gamma_\mathcal{B}$, using linear or binary (tree) search.
 2. Find an unused particle i compatible with species α or allocate storage for a new particle if none is found.
 3. Sample position \mathbf{r}_i from a problem-dependent probability density function, for example, uniformly inside the simulation box. Insert the particle into the corresponding cell list \mathcal{L}_c . If using NNs, build a new neighborhood \mathcal{N}_i .
 4. Set any additional attributes, such as particle size R_i , according to the inserted species α .
 5. Mark the particle as unprotected, $\mathcal{P}_i = \mathcal{C}_i$, and set $t_i \leftarrow t$.
 6. Schedule an immediate insertion event $t_e \leftarrow t$, $p = 0$, $\nu = -1$ and insert i into the event queue.
-

E. Particle Protection

For better performance, one can try to use the freedom afforded by the FPKMC algorithm to select particle protection so as to delay as far as possible the very next event in the queue. However, finding an optimal space partitioning is a difficult problem of non-linear optimization, especially since events other than the first-passage and no-passage propagations are taking place concurrently during the simulation. Here we present a practical algorithm that works well enough.

Our algorithm for protecting a given particle i finds the nearest pair of unprotected particles whose protection affects the protection of i . It is recursive and rather complex in its details, as described in detail in Algorithm 4. Here we try to give a more intuitive and brief verbal explanation, which can be used to design alternative implementations. The algorithm starts from the particle i and finds the maximal possible size of its protective region by examining all of the neighboring objects limiting the protection and finding the “nearest” (most-limiting) neighbor and also the “next-nearest” (next-limiting) neighbor. These neighbors could be other protected and unprotected particles, nearby hard walls, the cells used to build the LLCs, or the bounding neighborhood \mathcal{N}_i . For each of these cases one can calculate the maximal allowed size of the protection \mathcal{P}_i afforded by the neighboring object. Specifically, if the neighboring object is itself an unprotected particle it is assumed that the particles would be protected with touching protections whose sizes are proportional to the square roots of the diffusion coefficients. If the limiting neighbor is an unprotected particle, the algorithm recurses by repeating the process with that particle replacing particle i . The recursion continues until a neighbor is found whose own limiting neighbor is the particle i , that is, a pair of mutual nearest neighbors i and j is found. If i and j are sufficiently close and the next-nearest neighbors of i and j allow for pair protection (see Section III B), the particles i and j are protected as a pair. Otherwise, they are protected with touching single-particle protections. The recursion trail is then reversed and all the particles visited during the recursion are protected with the maximal allowed protection, accounting for the recent protection of their nearest neighbor and the previously-identified next-nearest neighbor. At the end of the process, the particle i and possibly a number of other particles are protected and the neighbor limiting the size of \mathcal{P}_i has been identified.

Algorithm 4: Protect the unprotected particle i with the largest possible \mathcal{P}_i and report

the neighbor l_{touch} that is limiting further growth. The radius of the protective region should not exceed an upper bound R_{max} on the largest possible size of any protection in the system. Note that the same algorithm applies to either spheres or cubes if distances are measured in the appropriate metric.

1. Set $l_{touch} \leftarrow 0$. If particle i is immobile, return without protecting it.
2. Push the current particle under consideration i , its limiting neighbor $l = 0$, and the maximum allowed protection size due to that neighbor $R_{max}^l = \infty$ on the recursion stack.
3. Read i , l_{old} and R_{max}^{old} from the top of the stack. Add room on the stack to store the nearest neighbor l and the associated $R_{max}^{\mathcal{P}}$ and, similarly, \tilde{l} and $\tilde{R}_{max}^{\mathcal{P}}$ for the second-nearest neighbor and marker $I_l^{\mathcal{P}}$ indicating if l was protected when it was first encountered.
4. Select the neighbor search range $R_{search} > R_i + R_{max}$ such that all particles closer than $R_i + R_{max}$ are inspected as potentially limiting. Set $R_{max}^{\mathcal{P}} \leftarrow \min(R_{search}/2, R_{search} - R_{max}) > R_i$ to guarantee that \mathcal{P}_i is not too large, $R_i^{(\mathcal{P})} \leq R_{max}^{\mathcal{P}}$, i.e., it does not overlap with any protections not inspected during the neighbor search.
5. Set $\tilde{l} \leftarrow 0$ and $\tilde{R}_{max}^{\mathcal{P}} \leftarrow R_{max}^{\mathcal{P}}$. If $R_{max}^{old} > R_{max}^{\mathcal{P}}$ set $l \leftarrow 0$, otherwise, set $l \leftarrow l_{old}$, $R_{max}^{\mathcal{P}} \leftarrow R_{max}^l$.
6. Iterate over all the particles or hard walls j closer than R_{search} (if>NNLs are used, include $j = \mathcal{N}_i$ in the search to ensure that $\mathcal{P}_i \subset \mathcal{N}_i$):
 - (a) If i and j are of non-interacting species or if $j = l$, skip this neighbor.
 - (b) Calculate the distance R_{ij} between i and j .
 - (c) If j is protected, set $R_{\mathcal{P}} = R_{ij} - R_j^{(\mathcal{P})}$, otherwise set $R_{\mathcal{P}} < R_{ij}$ so that R_{ij} is partitioned among the two particles according to their diffusion coefficients (see Section III B).
 - (d) If $R_{\mathcal{P}} < R_{max}^{\mathcal{P}}$ or $R_{\mathcal{P}} < \tilde{R}_{max}^{\mathcal{P}}$, update l , $R_{max}^{\mathcal{P}}$, \tilde{l} and $\tilde{R}_{max}^{\mathcal{P}}$ accordingly.
7. Continue to the next or end the current level of the recursion (i.e., pop the top of the execution stack and if the stack depth is positive resume from the previous step 7e). If l is a particle then set $I_l^{\mathcal{P}} \leftarrow \mathcal{T}$ if it is protected, and $I_l^{\mathcal{P}} \leftarrow \mathcal{F}$ otherwise:

- (a) If $l = \mathcal{N}_i$ then set $R_j^{(\mathcal{P})} \leftarrow R_{max}^{\mathcal{P}}$. If $R_j^{(\mathcal{P})} \leq (1 + \delta_{min})R_i$ set $l_{touch} \leftarrow \mathcal{N}_i$ and return from the recursion.
 - (b) Else if $l = -w < 0$ is a hard-wall then partner the particle with the wall if possible (we omit the details for brevity) and return from the recursion.
 - (c) If $l = l_{old} > 0$ then i and l are nearest neighbors. Store $R_{j,max}^{\mathcal{P}} = \tilde{R}_{max}^{\mathcal{P}}$ and $l_j = \tilde{l}$ and return from the recursion.
 - (d) Else if $l = j > 0$ is an immobile particle, set $l_j = 0$, $\tilde{R}_{max}^{\mathcal{P}} = R_j$ and return from the recursion.
 - (e) Else if $l = j > 0$ is an unprotected particle, push $i \leftarrow j$, $l \leftarrow i$ and $R_{max}^l \leftarrow R_{ij} - R_{max}^{\mathcal{P}}$ on the stack, and continue the recursion by going back to step 3.
8. If $j = l > 0$ is an unprotected particle, set $R_{i,max}^{\mathcal{P}} = \tilde{R}_{max}^{\mathcal{P}}$ and $l_i = \tilde{l}$, and test whether nearest neighbors i and j can be protected as a pair:
- (a) If $R_{i,max}^{\mathcal{P}} \leq (1 + \delta_{pair})R_i$ and $R_{j,max}^{\mathcal{P}} \leq (1 + \delta_{pair})R_j$, determine the smallest appropriate protection sizes $R_{i,min}^{\mathcal{P}}$ and $R_{j,min}^{\mathcal{P}}$ that allow a meaningful pair protection (see Section IIIB). If $R_{i,max}^{\mathcal{P}} \geq R_{i,min}^{\mathcal{P}}$ and $R_{j,max}^{\mathcal{P}} \geq R_{j,min}^{\mathcal{P}}$, protect ij as a pair:
 - i. Choose $R_i^{(\mathcal{P})}$ appropriately, $R_{i,min}^{\mathcal{P}} \leq R_i^{(\mathcal{P})} \leq R_{i,max}^{\mathcal{P}}$, and similarly for j if it is mobile. Update R_{max} to be larger than both $R_i^{(\mathcal{P})}$ and $R_j^{(\mathcal{P})}$.
 - ii. Set $p_i \leftarrow j$ and set $p_j \leftarrow i$ if j is mobile.
 - iii. Let $k = \min(i, j)$ and $k' = \max(i, j)$ if j is mobile, otherwise let $k = i$ and $k' = j$. Set $\nu_k \leftarrow 0$, $t_k^e \leftarrow 0$ and delete k' from the event queue.
 - iv. Return from the recursion.
 - (b) Set $R_i^{(\mathcal{P})} \leftarrow R_{max}^{\mathcal{P}}$, update $R_{max} = \max(R_{max}, R_i^{(\mathcal{P})})$, and similarly for j if it is mobile. If $R_{max}^{\mathcal{P}} \leq (1 + \delta_{min})R_i$ then set l_{touch} to be the limiting (closer) neighbor of l_i and l_j , and return from the recursion.
9. Else if $j = l > 0$ is a protected particle:
- (a) If $I_l^{\mathcal{P}} = \mathcal{T}$ then set $R_i^{(\mathcal{P})} \leftarrow R_{max}^{\mathcal{P}}$. If $R_{max}^{\mathcal{P}} \leq (1 + \delta_{min})R_i$, report this neighbor as limiting, $l_{touch} \leftarrow l$.

- (b) Otherwise re-calculate R_{ij} from $R_{max}^{\mathcal{P}}$ and set $R_i^{(\mathcal{P})} \leftarrow \max(\tilde{R}_{max}^{\mathcal{P}}, R_{ij} - R_j^{(\mathcal{P})}) \geq R_{max}^{\mathcal{P}}$ to ensure the largest possible protection.
-

F. Other Steps in the Algorithm

There are several other common steps in the FPKMC algorithm that are implemented as separate subroutines in our Fortran 95 code for which we do not give detailed pseudocodes but list them below with brief notes.

- *Processing a collision between two particles.* This step is very application-specific because of the different types of reactions that may occur. Typically the processing involves deleting some particles and then possibly inserting others.
- *Processing a decay reaction.* This is also application-specific and consists of deleting the decaying particle and then inserting the reaction products at desired positions.
- *Un-protecting particle i .* The partner is identified and the appropriate probability distribution is used to sample a new position for the particle and also its partner if the partner is mobile. The LLCs are updated accordingly, \mathcal{P}_i and possibly \mathcal{P}_j are destroyed, and immediate protection events scheduled for particle i and its partner j if it is mobile.
- *Un-protecting all partners of an immobile particle.* This is necessary when the state of an immobile particle changes (e.g., it decays). First, a neighbor search identifies the partners of the immobile particle, they are unprotected and immediately scheduled for re-protection. Note that the neighbor search here relies on the LLCs and it is not necessarily safe to modify LLCs until the search completes.
- *Scheduling and processing of pair events.* Since implementation depends on the particle shapes and the types of reactions considered, we do not give detailed pseudocodes for this step. Section III discusses pair propagators for hard spheres.
- *Resetting the time counter to $t = 0$.* This step is useful for minimizing round-off errors, especially before an event generating other events with very small timestamps occurs. For example, insertions of cascades of defects creates dense lumps of particles

that evolve at time scales Δt comparable to numerical precision (10^{-16}) relative to the time scales of the majority of events, and thus $t + \Delta t \approx t$ due to round-off. This can be avoided by setting $t = 0$ after subtracting the current t from all time counters, including the particle times t_i and the event predictions t_i^e .

- *Synchronizing all particles.* It is occasionally useful to bring the whole system to the current point in time for analysis, saving the configuration to a file, etc. This requires un-protecting all particles. This is a good occasion to reset current time $t \leftarrow 0$ to avoid the round-off problems.

G. Optimizing Runtime Parameters for Efficiency

The most important parameter defining the performance of the FPKMC algorithm is the size of the cells used for neighbor searches. The common wisdom for a homogeneous system of identical particles is that, optimally, there should be about one particle per cell which balances the cost of neighbor searches with the cost of updating the LLCs and moving particles between the cells. However, this prescription does not necessarily apply to the often heterogeneous (in both time and space) systems encountered, for example, in radiation damage simulations. Some clusters of defects can grow to sizes more than a decade larger than monomer defects. Additionally, the system's evolution can entail disparate timescales differing by many orders of magnitude, from fast relaxation on the scale of picoseconds during initial cascade insertion to slow annealing on the scale of years.

For simplicity, the following discussion focuses on LLCs, without>NNLs. Furthermore, we assume that among several particle species present in the system there is a single highly mobile species α_m , such as interstitials in radiation damage modeling. In such cases, the bulk of computational effort is spent on protecting and propagating particles of the fast species. Generally, it is a good idea to assign as large protection as possible for the fast particle(s) without, however, having to search for too many neighbors while building a protective region. In our algorithm the protection size is limited by the smaller of the following:

1. The range for the neighbor search $R_s = \min(n_s L_c / 2, n_s L_c - R_{max})$, where $n_s = \lceil (R_{max} + R_p) / L_c \rceil$ is the number of neighboring cells to be searched in each direction, L_c is the linear dimension of the (cubic) cells, R_p is the radius of particle species

α_m , and $R_{max} \gg R_p$ is the current maximum size of protection in the entire system (typically an immobile cluster).

2. The distance R_v to nearby neighbors against which pair protection is not possible (due to third particles blocking it or the neighbor being too far away to make pair protection advantageous). This is a measure of the void size around the fastest particles.

The optimal performance is achieved when the two bounds are approximately equal, $R_s \approx R_v$, i.e., the number of searched neighbors is just enough to find the largest possible protection, no more, no less. The void size R_v can be measured for a given configuration by gradually enlarging the cells until the average size of protective domains stops increasing, becoming close to R_v . Numerical tests have confirmed that indeed the choice of L_c such that $n_s = 1$ and $R_s \approx R_v$ is optimal.

In an actual simulation it is too expensive to estimate R_v at every step but it is still possible to use an adaptive method for selecting a cell size. Specifically, we monitor how many protections for particles of species α_m have been limited by cell size (i.e., by R_s) and how many have been limited by nearby particles (i.e., R_v). We observe that it is best to keep the former a small but nonzero fraction of the latter. If the runtime statistics show that too many protections are blocked by the cell size, the cells are enlarged by reducing the number of cells by one along each dimension of the simulation volume. Conversely, the cell count is increased by one if the statistics show that too few protections are blocked by the cells. The cell size strongly affects the performance of the code. If too small, the protections will be small too leading to shorter scheduled propagation times and, thus, slower time evolution. If the cells grow too large, there will be many neighbors to examine during each protection, slowing down the calculations. The use of NNs and, in particular, BSCs (see Section IV A 2), becomes advantageous when very large clusters are present, so that $n_s = 1$, i.e., $L_c > R_{max} + R_p$ and the cells contain many smaller particles at sufficiently high monomer densities. In the simulations of irradiated materials reported in section V B, we focus on low particle densities and find that the use of NNs is not necessary to achieve an optimal performance.

Note that computational cost is not always dominated by a single ultra-fast species. As an example, consider the case of modeling radiation damage inflicted in the form of defect pairs consisting of a very mobile interstitial and a much less mobile vacancy. At

first, the interstitial propagation events will dominate the event loop and the interstitials will quickly diffuse to absorbing sinks, such as hard-wall boundaries or nearby vacancies, and disappear. This will leave behind the slower vacancies that will continue their random motion until the next defect pair is inserted. If the insertion rate is low, the vacancies can move significantly between successive insertions and after the fast interstitials all died out. During such intervals the computational cost is dominated by vacancy propagations and protections. In such conditions, the focus of particle protection will have to shift from interstitials to vacancies and back to interstitials. In general, the choice of optimal cells and protection sizes is complex and problem-dependent. Our implementation of the FPKMC algorithm collects various statistics that can be used to make runtime adjustments and improve the simulation performance.

H. Mixing Time-Driven with Event-Driven propagations

Under certain conditions the exact event-driven handling of particle diffusion may become inefficient and/or cumbersome. For example, in a dense group of closely spaced particles, protection of single particles and pairs is severely limited by the third particles in close proximity. In such conditions, particle displacements can become too small to deserve asynchronous event-driven handling. Time-stepping avoids the cost of event queue operations and simplifies overlap detection. Therefore, it can be more efficient to use time-stepping for dense groups of particles, similar to the time-driven hopping algorithm presented in Section III B 2 that avoids the use of complex pair propagators for spheres. Use of small hops is also advantageous when particles or surfaces (e.g., grain boundaries) have complex shapes making analytical treatment of particle diffusion and collisions difficult or impossible. Yet another example when simple time stepping is useful is tightly-bound collections of particles (clusters) that may act as a single particle and have complex internal structure and dynamics (relaxation). Adding a time-driven component to the asynchronous event-driven FPKMC algorithm allows to retain the overall algorithm efficiency even in such difficult conditions.

The particles which cannot be protected by a sufficiently large protection are marked as *time-driven particles* and are *not* inserted into the event queue. A special type of event, a *time step event*, is introduced and always scheduled to occur at equal time intervals Δt . When a time-step event is processed, all time-driven particles are moved simultaneously,

followed by overlap checks and reactions, if any detected. Particles that are time-driven are not protected against each other, instead, they can stay unprotected or are protected only against the event-driven particles. Time-driven particles whose protections overlap form a cluster and are propagated synchronously with the same time step Δt tailored to the fastest particle in the cluster. For simplicity, all time-driven particles may be treated as one cluster with a single global Δt . However, it is often the case that different species have widely differing diffusion coefficients and therefore very different time-steps will be appropriate for different species. To solve this problem, one can use the n -fold (BKL) [15] synchronous event-driven algorithm inside each cluster and replace the time step events with *BKL hop events*. At each hop event all particles of a single species move by a small but non-negligible distance while all other particles remain in place. This way, the more mobile particles move more frequently (with correct relative frequencies) than the less mobile ones.

Note that the particles in such a time-driven cluster can take hops up to the time of the next event in the queue, since it is known that the hops can not be preempted by another event. In some situations this may improve efficiency by reducing the number of heap operations and also increasing the locality of the code by focusing multiple events on the same (cached) small group of particles. We are currently developing an implementation of a mixed event-driven (first-passage) with time-driven (n -fold KMC) algorithm and will report additional details and results in a future publication.

V. VALIDATION AND RESULTS

In this section we apply the FPKMC algorithm to several diffusion-reaction problems of increasing complexity. To validate the new algorithm and to demonstrate its efficiency, we compare our simulations to results obtained using two different object KMC (OKMC) codes developed earlier for simulations of continuum diffusion (BIGMAC code [1]) and for simulations of random walks on a lattice (LAKIMOCA code [6]). Presented in Section V A, the first two test problems are relatively simple validation studies for the case of two-species annihilation, $A + B \rightarrow 0$, when the two species have different diffusion coefficients. FPKMC simulations for this model are compared against results obtained with the BIGMAC code. In Section V B 1, we apply FPKMC to a more challenging test problem of damage accumulation in a metal thin film subjected to electron irradiation. We base our FPKMC

simulations on a well known model of α -iron studied earlier using three different methods, including cluster dynamics, lattice OKMC, and the approximate continuum OKMC algorithm JERK Ref. ([5]). Here we compare our results for this model against simulations performed using the LAKIMOCA lattice-based code. Finally, in Section V B 3 we apply FPKMC to simulations of radiation damage accumulation over previously inaccessible time scales, namely, to time intervals and radiation doses characteristic of material lifetime in a nuclear reactor. To our knowledge, this is the first time an atomistic model has reached technologically relevant radiation doses exceeding the previous simulation benchmarks by several orders of magnitude.

A. Two-Species Annihilation

As a validation study, let us first consider a system of spherical particles of two species A or B in three dimensions. The particles of different species have different radii R_A and R_B and different diffusion coefficients D_A and D_B , $D_B < D_A$. Particles of like species do not see each other but particles of unlike species annihilate upon hard-sphere contact, i.e. at the annihilation distance $r_{A+B \rightarrow 0} = R_A + R_B$. As a base for comparison, we first simulated the same model reaction using the BIGMAC code for several values of the hop distance Δ , characterized hereafter by the dimensionless ratio $\delta = \Delta/(R_A + R_B)$. For the hopping-based pair propagators in FPKMC we set $\delta = 0.1$, which was found to be sufficiently small to give accurate results, yet large enough to make the pair propagators almost as efficient as the analytical pair propagators for cube-shaped particles.

The simulations were performed in a cubic simulation domain of volume L^3 with periodic boundary conditions. In simulations reported here, half of the particles are of species A and the other half are of species B . We consider two different initial conditions. In the first case A s and B s are randomly and uniformly distributed in the simulation volume (the overlapping particles of unlike species are removed). The reaction kinetics is described by the reduction of the number of A (or B) particles with time, $N_A(t) = N_B(t)$. Figure 2 shows this decay kinetics for FPKMC simulations as well as for BIGMAC runs with different hop sizes δ . The results indicate that hop sizes as large as $\delta = 1/4$ can be used in BIGMAC without a noticeable error and that BIGMAC and FPKMC produce virtually indistinguishable results.

In the second test case, one half of the box ($x \geq L/2$) is randomly filled with A 's and

the other half ($x < L/2$) with B 's. The resulting $N_A(t) = N_B(t)$ is shown in Fig. 2, with similar qualitative behavior as in the first case of intermixed A s and B s. The concentration profiles $c_A(x; t)$ and $c_B(x; t)$ are shown for several different points in time in Fig. 2 for both FPKMC and BIGMAC with hop size $\delta = 1/4$ revealing an excellent agreement between the two algorithms. The efficiency of BIGMAC simulations is proportional to the square of the hop size, but even for the rather large hop sizes used in our BIGMAC simulations, the latter take several hours to annihilate most of the initial $1.28 \cdot 10^5$ particles, whereas the FPKMC code accomplishes the same in about 10 minutes on a typical single-processor desktop machine essentially independently of the value of δ .

B. Radiation Damage

Modeling of radiation damage in reactor materials is not only another test problem for our algorithm but also a technologically important application, notably in the design and maintenance of nuclear power plants. When a metal is irradiated, incoming high energy particles (neutrons, ions, electrons) collide with atoms of the host crystal lattice inducing displacement cascades and producing numerous defects, such as excess vacancies and interstitials. Many of these defects quickly annihilate with each other, but some diffuse away from the initial impact locations eventually finding other defects to react with and to form defect clusters. The cluster density and sizes can grow over time resulting in substantial (most often detrimental) modifications of material microstructure and properties. Atomistic KMC method is well suited for simulations of radiation damage accumulation by tracing the numerous diffusive hops and reactions among crystal defects induced by collision cascades. Unfortunately, the same detailed nature of KMC simulations makes them computationally demanding and limits their time horizon to times far shorter than the technologically relevant time scales (years).

The FPKMC algorithm seems to be perfect for simulations of materials under irradiation and, in fact, it was this particular application that served as an initial inspiration for the new method development. The advantage of FPKMC over other KMC methods is that the new algorithm handles equally efficiently both the fast stages of cascade annealing when the density of diffusing defects is high and the slow evolution of the small number of surviving defects that continue their diffusive motion during the intervals between subsequent collision

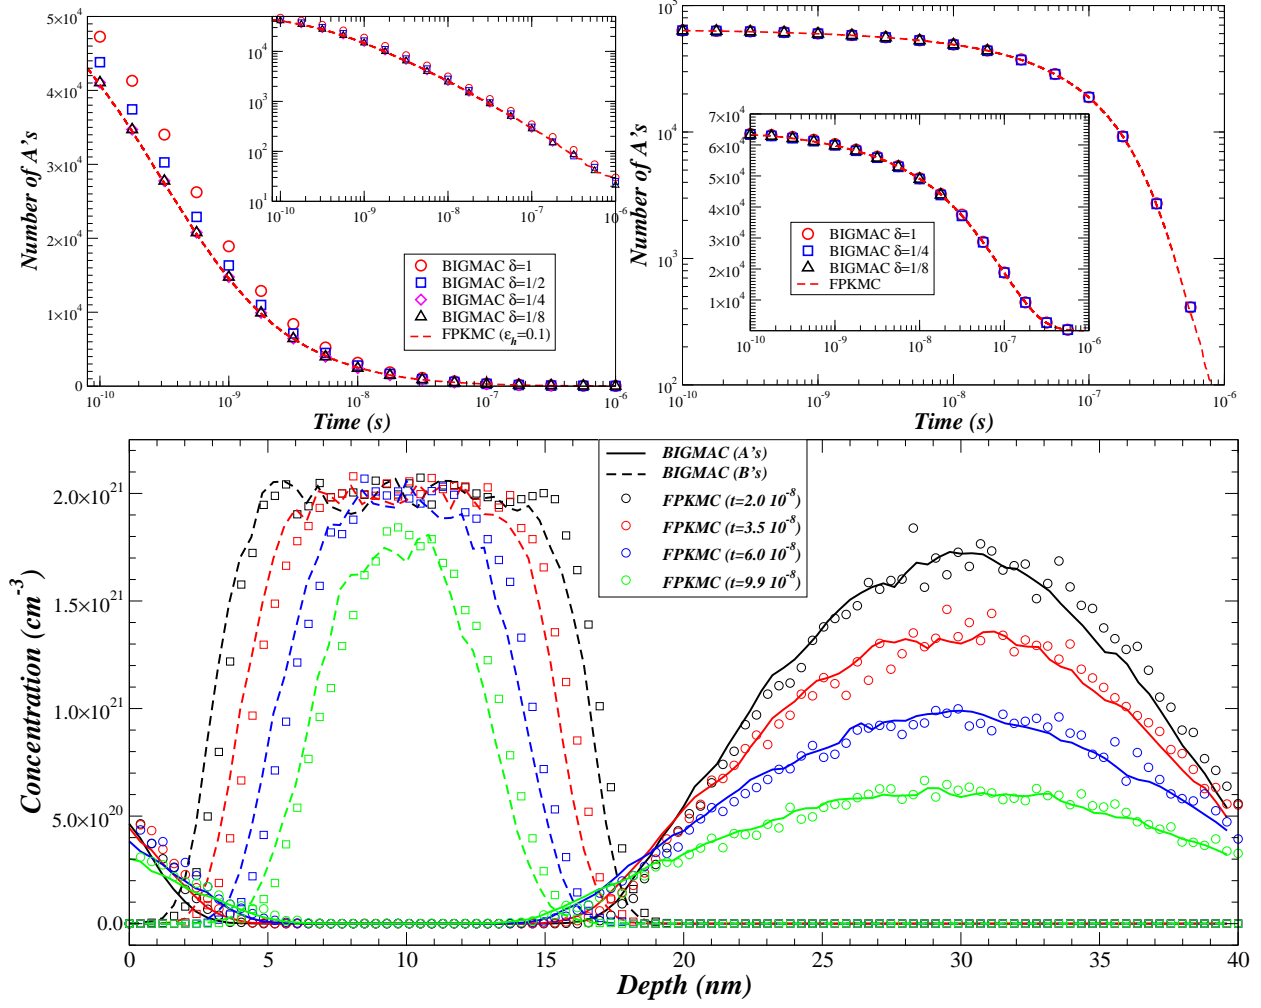


Figure 2: Comparison between BIGMAC and FPKMC for the two-species annihilation ($A+B \rightarrow 0$) test problem in three dimensions, starting from a uniformly-mixed (top left) and phase-separated (top right and bottom) initial state. Two top panels and their insets show the number of particles as a function on time, $N_A(t) = N_B(t)$, on log-log and semi-log scales. The bottom panel shows the particle concentration profiles $c_A(x;t)$ (solid curves for BIGMAC, circles for FPKMC) and $c_B(x;t)$ (dashed curves for BIGMAC, squares for FPKMC) for the phase-separated initial conditions. Concentration profiles taken at several times during the simulation are plotted in different colors.

cascades. This intermittent (fast-slow-fast-...) character of system's evolution is the major computational challenge in radiation damage simulations that the asynchronous FPKMC algorithm addresses.

To enable radiation damage simulations, several particle species are introduced:

monomers, including highly-mobile interstitials (I) and less-mobile vacancies (V), a number of mobile *cluster species*, for example, dimers (I_2 and V_2) and trimers (I_3 and V_3), immobile species representing clusters larger than any of the mobile species (I_c and V_c), and defect (Frenkel) pairs (IV). Each particle is assigned a hard-sphere radius: for the clusters the radius is related to the number of monomers $c \geq 1$ contained in the clusters assuming that the monomer volumes are additive, i.e. $R_c \sim R_0 + (R_1 - R_0)c^{1/3}$.

Frenkel pairs (IV) are inserted in the simulation volume at a specified birth rate and instantly (i.e., with decay time $\tau_{IV} = 0$) decay via $IV \rightarrow I + V$. The resulting interstitial and vacancy monomers are placed randomly within the simulation box, either at some initial distance from each other or completely independently of each other. Electron irradiation creates individual Frenkel pairs while irradiation by high-energy ions or neutrons creates damage in the form of *displacement cascades* producing compact collections of monomers and clusters, each cascade containing about 100 Frenkel pairs. The collision cascades are randomly selected from a library of cascade configurations generated in Molecular Dynamics simulations and the cascade locations and orientations are sampled from appropriate distributions.

Upon collisions, particles of like species coalesce, for example, $I + I \rightarrow I_2$ or $V + V_3 \rightarrow V_{c=4}$, whereas collisions of particles of unlike species lead to complete or partial annihilation, for example, $I_2 + V_{c=4} \rightarrow V_2$. In our current implementation, the distance at which two particles collide must be equal to the sum of their radii since the FPKMC algorithm handles the geometry of the protective regions assuming that an additive hard-sphere interactions among the particles. For consistency, the same is assumed in simulations performed with the lattice-based LAKIMOCA code that are used here for comparison. To imitate the stronger effect of elastic strain on the interstitials (the bias), the interstitials are assigned radius 1.2 times larger than that of the vacancies.

The defect clusters emit monomers at a given rate, represented as a decay reaction, for example, $V_{c=5} \rightarrow V_{c=4} + V$, or $I_2 \rightarrow I + I$. The emitted monomer is placed at a preset *emission distance* δ_e from the cluster surface, in a random direction. This is at variance with most other commonly adopted non-local emission rules, e.g. in the mean-field Cluster Rate Theory method [21], in which the monomers are emitted from the clusters to “infinity”. If and when a cluster shrinks by emission to a size at which it becomes mobile, the species of the remaining cluster is changed accordingly, for example, $V_{c=4} \rightarrow V_3 + V$.

The rates of defect diffusion and monomer emission from the clusters are calculated according to the standard expressions for the rates of thermally activated rate processes in solids. Defect migration and binding energies needed to compute the rates are calculated atomistically, e.g. from first-principles theory, or estimated from experimental data. In the simple model we consider here, only the monomers are mobile, with diffusion coefficient

$$D_1 = D_0 e^{-E_m/kT},$$

where E_m is the activation energy for defect migration (lattice hop). The rate of emission of monomers from a cluster composed of c monomers is

$$\Gamma_c = \Gamma_0 D_1 a^{-2} c^{2/3} e^{-E_b(c)/kT},$$

where a is the lattice spacing, Γ_0 is a constant that depends on the lattice type and $E_b(c)$ is the monomer binding energy in a cluster of size c , estimated using

$$E_b(c) = E_f + [E_b(2) - E_f] \frac{c^{2/3} - (c-1)^{2/3}}{2^{2/3} - 1},$$

where E_f is the monomer formation energy [5].

1. Thin film of metal under electron irradiation

As a test problem, here we consider a model previously studied in Ref. [5] using two other KMC algorithms and the mean-field cluster dynamics [21]. The model system is a $0.287\mu m$ -thick film of α -iron subjected to electron radiation. Periodic boundary conditions are used in the x and y directions, while absorbing walls are used in the z direction. Further details of the model are given in Ref. [5] and will not be repeated here. We made a few minor changes to the model parameters presented in Ref. [5], notably switching from the non-local to local emission of monomers from the clusters. Here we compare our FPKMC simulations with new results obtained using the LAKIMOCA code [6].

In pure iron, the interstitials are much more mobile than the vacancies which results in their rapid absorption at the free surfaces or annihilation with the vacancies. Due to this very high monomer mobility, very few if any interstitial clusters form. However, the ones that do nucleate, can grow rather large because the binding energy of interstitials in the clusters is rather high. A large fraction of computational effort is therefore expended

on propagating the interstitials. However, depending on the irradiation rate (fluence), the vacancies may have time to move significantly between successive insertions of Frenkel pairs. Here we present a high-fluence case studied computationally and experimentally in Ref. [5]. Further computational tests of the FPKMC code in a wide range of irradiation conditions will be discussed in Section VB 2.

A comparison of the simulated time evolution for the number of defects, as well as their spatial distribution at the end of 120 seconds of irradiation, is shown in Fig. 3. Generally, we observe a good agreement between FPKMC and LAKIMOCA results although the number of accumulated defects is slightly smaller in the FPKMC simulations than in the LAKIMOCA runs. This is perhaps a consequence of different treatment of defect diffusion in the two methods: in FPKMC the defects move by continuous diffusion whereas in LAKIMOCA the defects walk on a lattice. However, at sufficiently low densities, lattice discreteness is not as important and the effective reaction and emission rates are well matched between the discrete and continuum models. Further comparisons between FPKMC and OKMC simulations will be given in a future publication, here we simply observe the ability of the FPKMC algorithm to correctly simulate radiation damage by comparing it to a standard OKMC. Each FPKMC simulation sample used for obtaining the data plotted in Fig. 3 has taken less than five minutes to complete on a modest workstation (3.4GHz Xeon 64-bit processor). Efficiency of FPKMC simulations is further discussed in the next section.

2. Performance

For a large number of particles N , computational complexity of the FPKMC algorithm per event should be order $O(N \log N)$ which is the cost of event queue updates. For a typical value of $N = 10^5$, the logarithm is masked by other dominant costs that are all linear (e.g., neighbor searches and sampling from the propagators). Computational tests in the range of N from 10^4 to 10^6 have indeed verified that the cost of FPKMC simulations scales linearly with the number of particles.

The FPKMC code performance also depends on a number of other parameters, e.g. particle density, the disparity of diffusion and emission rates, differences in particles sizes between the different species, etc. These and other factors and their interactions affect the overall performance in complex ways that are yet to be fully examined. We defer to future appli-




Figure 3: Comparison between FPKMC (symbols) and LAKIMOCA [6] (lines) simulations of a $0.287\mu\text{m}$ -thick film of α -iron subjected to 120 seconds of electron radiation at temperature $T = 200^\circ\text{C}$. The results shown on the plots are obtained by averaging over 50 runs. (Top) The time evolution of the total number of mono-vacancies, small vacancy clusters and interstitials (in clusters). The statistical error bars are comparable to the symbol sizes. The inset shows a different scale to focus on the smaller clusters. (Bottom) The density profile along the thickness of the film for all vacancies (in monomers and clusters), mono-vacancies, vacancies in clusters and all interstitials (in monomers and clusters) at the end of 120 seconds of simulated irradiation. The error bars are comparable to the symbol sizes, except for the interstitials for which the statistics is poor.

Dose rate (dpa/s)	Total dose (dpa)	Simulated time	Speed (s/cpus)	Efficiency (dpa/cpuday)
$1.5 \cdot 10^{-4}$	18	33 hrs	0.14	1.8
$1.5 \cdot 10^{-5}$	2.9	54 hrs	1.3	1.7
$1.5 \cdot 10^{-6}$	4.1	31 days	13	1.6
$1.5 \cdot 10^{-7}$	1.6	125 days	150	2.0
$1.5 \cdot 10^{-8}$	10	21 years	$2.1 \cdot 10^3$	2.7
$1.5 \cdot 10^{-9}$	8.4	175 years	$2.3 \cdot 10^4$	3.0

Table I: Performance of the FPKMC algorithm in simulations of a $0.287\mu\text{m}$ -thick film of α -iron at $T = 473^\circ\text{K}$ subjected to different fluxes (dose rates) of electron radiation. Conditions typical of nuclear reactors correspond to dose rates on the order of 10^{-8}dpa/s and lifetimes from years to several decades, while in the accelerated irradiation facilities the dose rates can be on the order of 10^{-4}dpa/s and the tests can last for several hours or days. A displacement dose of K -dpa means that, on average, each atom has been displaced from its equilibrium lattice position K times due to incoming radiation. Note that simulation efficiency changes over the course of one simulation so that the overall efficiencies reported in the last column are representative averages.

cations to study the subtle effects of various model parameters on the method performance. Here we present a few figures on the performance of our FPKMC code for the same simple model of metal thin film described in Section V B, as a function of irradiation flux measured in the units of displacements per atom (*dpa*) per second (*dpa/s*). Kinetics of defect microstructure evolution and, in particular, the amount and character of accumulated damage depend sensitively on flux and temperature. For example, the defect density (accumulated damage) increases with the increasing flux and/or the decreasing temperature. It is therefore not *a priori* obvious that FPKMC will be equally effective in dealing with a wide range of fluxes and temperatures. A reasonable measure of the algorithm performance is the damage dose simulated over a unit of CPU time expressed for example, in the unit of dose simulated in one day of computing (dpa/cpuday).

The figures presented in Table I demonstrate that, with the use of optimization techniques discussed in Section IV G, the algorithm performance remains nearly constant across several decades of radiation flux. This highly desirable property derives from the ability of our asynchronous event-driven algorithm to deal with very large differences in event rates and

local densities. We observe that, in the course of a single simulation, the FPKMC algorithm self-adjusts to the current conditions of spatial and temporal heterogeneity without much intervention or parameter tuning.

3. Radiation Damage on Reactor Timescales

Future development of nuclear energy demands materials that can withstand harsh conditions of particle irradiation, high temperature, mechanical stress and active chemical agents over tens of years. The only fully reliable method to evaluate the potential of a candidate material is to subject it to conditions relevant for the future reactor designs. However, such an approach is not practical given that the relevant environment can be achieved only after the reactor is already built. Furthermore, even if an appropriate material testing facility were to exist, testing the candidate materials over the intervals of 50 or 100 years would not be practical.

The idea of *accelerated material testing* is to subject candidate materials to conditions even harsher than in the reactor but over shorter periods of times, e.g. a few hours or days, in the hope that the observed (accelerated) material degradation can be used as a predictor of the performance of the same material during its lifetime in a real reactor. The premise of accelerated materials testing is that materials theory and numerical simulations can provide a reliable connection between the accelerated tests and the material lifetime performance predictions. To serve this purpose, material simulations should meet two conditions. First, accurate material models need to be developed and validated against experimental measurements. Second, the simulation algorithms need to be efficient to enable computational predictions of materials performance under reactor conditions. The performance data presented in Table I suggests that FPKMC can meet this second challenge. Here we continue to focus on the thin-film model studied in Section V B and use our FPKMC code to simulate damage accumulation at two different dose rates, a *high* dose rate of $1.5 \cdot 10^{-4}$ dpa/s typical of accelerated experiments in material testing facilities such as JANNUS [22], and a *low* dose rate of $1.5 \cdot 10^{-8}$ dpa/s typical of the existing nuclear reactors. We ran both simulations to a total dose of 10 dpa, which required several CPU days per sample on a common workstation.

Direct comparison of two simulations performed at the same temperature $T = 200^\circ\text{C}$

revealed very different kinetics and end-of-dose damage, which is not surprising given that much more time is available for damage annealing (healing) at the slow (reactor) time scales. In order to enable scaling from high dose rates to low dose rates, it was proposed [23] to raise the temperature in the high dose rate irradiation test so as to preserve the ratio of the damage insertion rate to the rate of defect diffusion. Such scaling would be exact if there were only one evolution mechanism whose rate can be adjusted by changing the temperature. However, even in the simple model of α -iron considered here, there is a whole spectrum of mechanisms and associated rates with different temperature activation parameters. One can only hope that an approximate scaling can be achieved by adjusting the rate of just one dominant mechanism that controls the overall rate of damage accumulation. In an extensive series of numerical experiments we observed that, within the simple model considered here, the overall rate and character of damage evolution appears to be controlled by the ratio of vacancy diffusion to the irradiation dose rate. This is likely because, at all temperatures of interest here, the interstitials are much more mobile than the vacancies and disappear nearly instantly following the insertion of a Frenkel pair, leaving the less mobile vacancies to diffuse and cluster in the absence of interstitials.

Figure 4 compares the state of damage reached at $T = 262^\circ\text{C}$ and high dose rate, with that reached at $T = 130^\circ\text{C}$ at low dose rate. The kinetics of damage accumulation and the resulting populations of voids are similar to each other but the scaling is only approximate and some differences in the resulting microstructures are noticeable. Most visibly, a few large interstitial clusters form at the high dose rate where the time interval between successive Frenkel-pair insertions is comparable to the lifetime of inserted interstitials. On the contrary, at the lower dose rate the newly inserted interstitials disappear (at the free surfaces or through annihilation with the vacancies) well before the next Frenkel pair is inserted. Additionally, as seen from the histograms shown in the insert, the average size of the vacancy clusters is larger at the higher dose rate. The approximate agreement in the overall swelling supports the idea that it may be possible to compensate the enhanced rate of damage accumulation in accelerated tests by raising the test temperature so that the resulting damage is approximately the same as in the reactor but at a lower temperature. That the scaling is only approximate is less than surprising considering that multiple rate processes, each with its own temperature dependence, act together to produce the resulting damage kinetics.

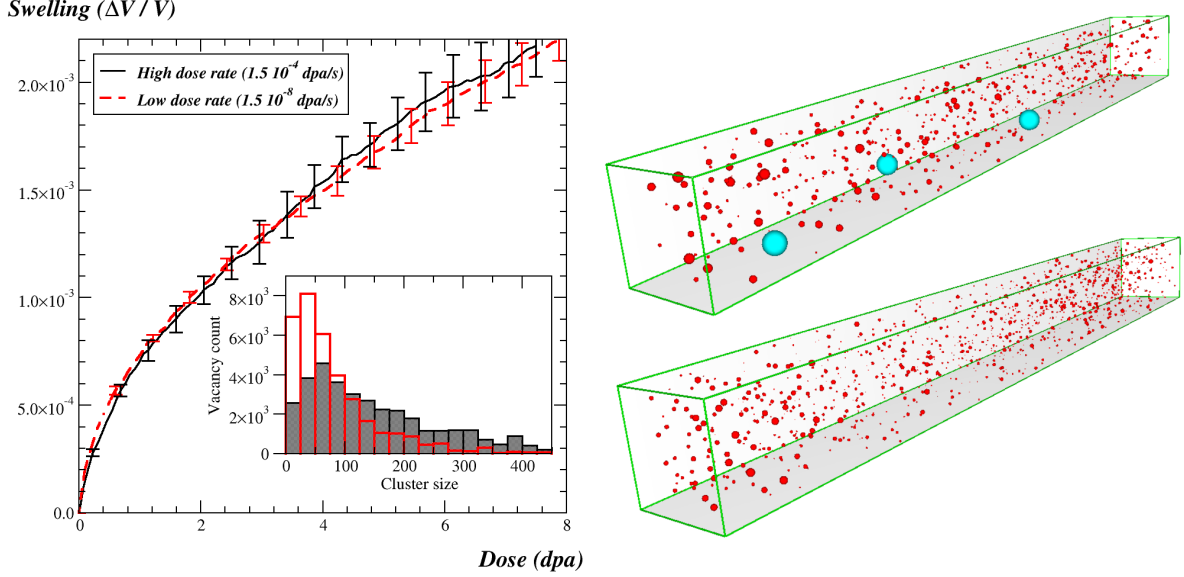


Figure 4: State of damage in a thin α -Fe film irradiated by electrons to the total dose of 10 dpa: (Top right) At the high dose rate of $1.5 \cdot 10^{-4}$ dpa/s and $T = 262^\circ\text{C}$ and (Bottom right) At the low dose rate of $1.5 \cdot 10^{-8}$ dpa/s and $T = 130^\circ\text{C}$.

(Left) Volume fraction of vacancies (swelling) as a function of damage dose. The vacancy fraction was obtained by counting together all vacancies in the simulation volume, both in the vacancy monomers or vacancy clusters, and dividing the sum by the total number of atomic sites. The black solid curve is the swelling kinetics under the high dose rate/high temperature and the dashed red curve is the same kinetics under the low dose rate/low temperature irradiation conditions. Both curves were obtained by averaging over 10 independent simulation runs for each irradiation condition, and the vertical bars show the estimated statistical errors. The inset in the bottom figure shows a histogram of the distribution of vacancy cluster sizes at a dose of 5 dpa for both the high dose rate/high temperature irradiation (shaded gray bars) and for the low dose rate/low temperature irradiation (red bars).

VI. CONCLUSIONS

In this series of two papers we have described two algorithmic implementations of an asynchronous event-driven algorithm for diffusion kinetic Monte Carlo simulations of diffusion-reaction particle systems, first proposed in Ref. [2]. The First Passage Kinetic Monte Carlo (FPKMC) algorithm avoids long sequences of small diffusive hops commonly used to bring

particles to collisions, by enabling large super-hops sampled from exact (semi)analytical solutions for diffusion Green’s functions (propagators) in spatially isolated protective regions each containing just one or two particles. The FPKMC algorithm is exact to the extent that the system’s stochastic trajectory is sampled from the exact diffusion-reaction Master Equation for the N -particle system, to within the accuracy of the single- and two-body Green’s functions. At the same time, for a range of simulations reported in the two papers, the new algorithm is several orders of magnitude more efficient than the existing (approximate) hopping-based algorithms.

In part I of this series we described the basic algorithmic ideas and elementary theory at the basis of the FPKMC method. The algorithm was implemented for cube-shaped particles, a simplifying case, and the ability of the new algorithm to overcome the bottleneck of low particle density was demonstrated. In part II (this paper) we extended the new method to considerably more complex simulations in which particle diffusion is just one among many competing rate processes. We described generalizations and algorithmic components necessary to handle systems with multiple particle species and multiple reactions, including annihilation, clustering, emission, and particle birth and death. We focused on the case of hard spheres that are more appropriate than the cubes for a multitude of intended applications. Handling hard sphere collisions as a first-passage process turned out to be more complicated than in the case of cubes which necessitated a hopping-based approximation for the pair propagators. The resulting solution to the sphere collision problem was found practical and efficient but raised a more general issue of combining the asynchronous event-driven framework employed in FPKMC with the more traditional synchronous time-driven simulations. Our successful implementation of time-driven pair propagators in FPKMC points to hybrid time-driven/event-driven algorithms as a promising direction for future research.

The new implementation of the FPKMC algorithm has proven suitable for simulations of damage accumulation in materials subjected to irradiation by high-energy particles. The accuracy of the FPKMC algorithm was validated on several test problems by comparing to traditional (object) KMC algorithms developed for diffusion-reaction simulations in the continuum (BIGMAC) and on the lattice (LAKIMOCA). The new algorithm is shown to perform well in a wide range of radiation conditions enabling, for the first time, simulations of irradiated materials to large technologically relevant radiation doses (e.g., 10dpa) on a serial workstation. Closing the gap between the short time horizon of traditional KMC simulations

and long material life in the reactor required to gain several orders of magnitude in computational efficiency. In FPKMC, this gain is achieved entirely through an exact factorization of the difficult N -body reaction-diffusion problem into one- and two-body problems.

With its efficiency, the new FPKMC method can make a significant impact on the important area of accelerated material testing for next-generation nuclear reactor designs. There is a strong synergy between atomistic KMC simulations and experiments to be carried out in accelerated testing facilities. The accuracy of an atomistic KMC model can be improved by expanding its mechanism catalog and obtaining more accurate values of model parameters. For this purpose, computationally efficient KMC simulations can be used to explore and identify experimental conditions in which accelerated material tests would be most informative for model validation, e.g. most sensitive to a particular mechanism or model parameter. Furthermore, the same simulations can be used to fine-tune the KMC models for conditions typical of real reactors. The approximate scaling predicted in our FPKMC simulations provides exactly the right kind of connection between simulations and experiments. Once the accuracy of the material models is established, FPKMC simulations can be used to extrapolate from accelerated material tests into relevant but inaccessible conditions of nuclear reactors without relying on any approximate scaling. To quantify reliability of such computational extrapolations, FPKMC simulations can be used to assess the uncertainties in computational predictions of accumulated damage given the uncertainties in model parameters, similarly to what is routinely done in climate modeling. Ultimately, efficient KMC simulations can and should become an integral component of reliable material testing programs.

In addition to irradiated materials used in nuclear power plants, semiconductor industry, space and medicine, the new method may prove enabling in other areas of sciences, engineering and finance. One particularly attractive application for the new method is in cell biology where multiple reaction-diffusion mechanisms conspire to produce a wide variety of specific biological responses. Parallelization of the FPKMC algorithm could further extend length and time scales accessible to diffusion-reaction simulations, even though the asynchronous event-driven algorithms are notoriously difficult to parallelize [12].

VII. ACKNOWLEDGMENTS

We would like to thank Christophe Domain, Pär Olsson, and Charlotte Becquart for helpful discussions and sharing with us their OKMC simulation data. We would also like to thank Mihai-Cosmin Marinica for testing our FPKMC code extensively and Maria-Jose Caturla for sharing with us her library of collision cascades. This work was performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was supported by the Office of Laboratory Directed Research at LLNL and the Office of Basic Energy Sciences U. S. Department of Energy.

-
- [1] S. K. Theiss, M.-J. Caturla, M. D. Johnson, J. Zhu, T. J. Lenosky, B. Sadigh, and T. Diaz de la Rubia. Atomic scale models of ion implantation and dopant diffusion in silico. *Thin Solid Films*, 365:219–230, 2000.
 - [2] T. Oppelstrup, V. V. Bulatov, G. H. Gilmer, M. H. Kalos, and B. Sadigh. First-Passage Monte Carlo Algorithm: Diffusion without All the Hops. *Phys. Rev. Lett.*, 97(23):230602, 2006.
 - [3] J. Dalla Torre, J.-L. Bocquet, N. V. Doan, E. Adam, and A. Barbu. JERK, an event-based Kinetic Monte Carlo model to predict microstructure evolution of materials under irradiation. *Philosophical Magazine*, 85:549–558, 2005.
 - [4] J. S. van Zon and P. R. ten Wolde. Green’s-function reaction dynamics: A particle-based approach for simulating biochemical networks in time and space. *J. Chem. Phys.*, 123(23):234910, 2005.
 - [5] A. Barbu, C. S. Becquart, J. L. Bocquet, J. Dalla Torre, and C. Domain. ”comparison between three complementary approaches to simulate large-fluence irradiation: application to electron irradiation of thin foils”. *Philosophical Magazine*, 85:541–547, 2005.
 - [6] C. Domain, C. S. Becquart, and L. Malerba. Simulation of radiation damage in Fe alloys: an object kinetic Monte Carlo approach. *Journal of Nuclear Materials*, 335(1):121–145, 2004.
 - [7] J. S. Reese, S. Raimondeau, and D. G. Vlachos. Monte Carlo Algorithms for Complex Surface Reaction Mechanisms: Efficiency and Accuracy. *J. Comp. Phys.*, 173(1):302–321, 2001.
 - [8] M. Biehl. Lattice gas models and Kinetic Monte Carlo simulations of epitaxial growth. In

- A. Voigt, editor, *Int. Series of Numerical Mathematics 149*, pages 3–18. Birkhaeuser, 2005.
- [9] M. Strobel, K.-H. Heinig, and W. Möller. Three-dimensional domain growth on the size scale of the capillary length: Effective growth exponent and comparative atomistic and mean-field simulations. *Phys. Rev. B*, 64(24):245422, Dec 2001.
 - [10] D. P. Tolle and N. Le Novere. Particle-based stochastic simulation in systems biology. *Current Bioinformatics*, 1(3):315–320, August 2006.
 - [11] S. J. Plimpton and A. Slepoy. Microbial cell modeling via reacting diffusive particles. *J. Phys.: Conf. Ser.*, 16:305–309, 2005.
 - [12] A. Donev. Asynchronous event-driven particle algorithms. *SIMULATION: Transactions of The Society for Modeling and Simulation International*, 85(4):229–242, 2008.
 - [13] B. J. Alder and T. E. Wainwright. Studies in molecular dynamics. I. General method. *J. Chem. Phys.*, 31:459–466, 1959.
 - [14] A. Scala, T. Voigtmann, and C. De Michele. Event-Driven Brownian Dynamics for Hard Spheres. *J. Chem. Phys*, 126(13):134109, 2007.
 - [15] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems. *J. Comput. Phys.*, 17(1):10–18, 1975.
 - [16] M. H. Kalos, D. Levesque, and L. Verlet. Helium at zero temperature with hard-sphere and other forces. *Phys. Rev. A*, 9(5):2178–2195, 1974.
 - [17] S. Redner. *A guide to first-passage processes*. Cambridge University Press, Cambridge, 2001.
 - [18] J. A. Given, J. B. Hubbard, and J. F. Douglas. A first-passage algorithm for the hydrodynamic friction and diffusion-limited reaction rate of macromolecules. *J. Chem. Phys*, 106(9):3761–3771, 1997.
 - [19] C.-O. Hwang, J. A. Given, and M. Mascagni. The Simulation-Tabulation Method for Classical Diffusion Monte Carlo. *J. Comp. Phys.*, 174:925–946, December 2001.
 - [20] A. Donev, S. Torquato, and F. H. Stillinger. Neighbor List Collision-Driven Molecular Dynamics Simulation for Nonspherical Particles: I. Algorithmic Details II. Applications to Ellipses and Ellipsoids. *J. Comp. Phys.*, 202(2):737–764, 765–793, 2005.
 - [21] A. Barbu and E. Clouet. Cluster dynamics modeling of materials: Advantages and limitations. *Solid State Phenomena*, 129:51, 2007.
 - [22] Y. Serruys, M.O. Ruault, P. Trocellier, S. Miro, A. Barbu, L. Boulanger, O. Kaïtasov, S. Henry, O. Leseigneur, P. Trouslard, et al. JANNUS: experimental validation at the scale of atomic

- modelling. *Comptes rendus-Physique*, 9(3-4):437–444, 2008.
- [23] M. J. Caturla, T. Diaz de la Rubia, and M. Fluss. Modeling microstructure evolution of fcc metals under irradiation in the presence of He. *Journal of Nuclear Materials*, 323(2-3):163–168, 2003.